ChipViz: Visualizing Memory Chip Test Data

Amit P. Sawant[a], Ravi Raina[b], and Christopher G. Healey[c]

[a,c]North Carolina State University, Department of Computer Science, Raleigh, NC, USA
[b]Qimonda AG, 3000 Centregreen Way, Cary, NC, USA

amit.sawant@ncsu.edu, ravi.raina@qimonda.com, healey@csc.ncsu.edu

Abstract. This paper presents a technique that allows test engineers to visually analyze and explore within memory chip test data. We represent the test results from a generation of chips along a traditional *2D grid* and a *spiral*. We also show correspondences in the test results across multiple generations of memory chips. We use simple geometric "glyphs" that vary their spatial placement, color, and texture properties to represent the critical attribute values of a test. When shown together, the glyphs form visual patterns that support exploration, facilitate discovery of data characteristics, relationships, and highlight trends and exceptions in the test data that are often difficult to identify with existing statistical tools.

1 Introduction

One of the biggest challenges in analyzing memory test data is discovering interrelationships between different test attributes. It is often time consuming and difficult to correctly interpret different test attributes using existing data analysis tools. With semiconductor manufacturing processes and technology changing rapidly, and test complexity increasing with every new generation of chips, its imperative that test data analysis tools keep pace.

The objective of memory testing (or, more generally, Integrated Circuit (IC) testing) is not only to isolate bad chips, but also to identify the root cause of failure which could be either weaknesses in chip design, or issues in manufacturing processes. Thus testing also acts as a feedback loop for design and manufacturing. With time to market the chip being absolutely critical for any IC company, its important that this feedback is integrated into the design and manufacturing process in a timely and efficient manner. This depends on how the data is organized, and more importantly, on how well it is presented for analysis and interpretation. One possible solution is to use visualization techniques to convert this large and complex dataset into a multi-dimensional visual image that the test engineers can use for exploring, discovering, comparing, validating, accounting, monitoring, identifying faults and process excursions, and studying the effects of adjusting different test parameters.

Numerous efforts have been made to use test data to improve yield and optimize tests. Previous research work includes developing a fault simulator to determine fault coverage of test patterns [1]. Sang-Chul et al. have developed an automatic failure analysis system based on production data [2]. Researchers have applied data-mining techniques to optimize VLSI testing [3]. Test visualization techniques have also been used

in area of software engineering to assist fault localization [4]. Recently Van de Goor et al. have developed methods to evaluate DRAM production test results to optimize tests and fault coverage [5].

The remainder of this paper proceeds as follows. In Section 2, we provide details on data collection. Section 3 describes our visualization technique. Sections 4 and 5 provide a few examples of visualizing memory chip test data. Finally, Section 6 discusses conclusions and future work.

2 Data Collection

The visualization process begins by collaborating with domain experts to identify the important parameters relating to yield and test optimization they want to analyze, explore, and monitor. The test datasets for this paper are taken from Qimonda AG[1], the fourth largest DRAM chip design and manufacturing company in the world. Generally memory test results are collected from manufacturing sites for analysis purpose. They consist of a spreadsheet report of the tests with corresponding failure rates and pertinent information about the individual test's attributes for each type of memory chip. It is then up to the test engineer to manually interpret the test results. The memory chip goes through different sets of tests called insertions. Each test typically has more than one critical attribute associated with it. Moreover, the same critical attributes may appear in multiple insertions. This leads to a complex dataset that is very large, and difficult to analyze and interpret correctly.

Data is viewed as a spreadsheet where rows represent individual tests and columns represent the attributes of the test. These attributes are typically related to critical timings and voltages of the chip, for example: *tRP*, *tRCD*, *tWR*, *tRAS*, *Vdd*, *Retention*, *Logistics*, *Current*, and *Failure Rate*. Attribute definitions are provided in the Appendix. The engineers currently depend on their experience and expertise to analyze the data and deduce meaningful information. Unfortunately, this is not an efficient method as the amount of data is huge and the interrelationships complex enough to confuse even the most experienced engineer.

The test datasets for this paper are taken from *DDR2/DDR3* memory chips. In this paper, we visualized the following four datasets:

1. ***Low Vdd at LT***: This dataset contains test results from a lot (a set of memory chips) with high failure rates at low temperature (LT) and low voltage (Vdd).
2. ***Retention at High Vdd and HT***: This case contains a certain memory chip with high retention failure rates on certain lots at high voltage and high temperature (HT).
3. ***tRCD and tRP at HT/LT***: This case contains high row to column address delay (*tRCD*) and row precharge time (*tRP*) failure rates across multiple critical attributes at high or low temperatures.
4. ***Optimized Test Data***: This dataset represents a stable, high volume product in which the tests/processes are already highly optimized and exceed the required pass thresholds.

[1] `http://www.qimonda.com`, formerly the memory chip division of Infineon Technologies AG

3 Visualization Technique

To visualize the memory test results, we adopted the following design guidelines proposed by Eick [6]: (1) ensure that the visualization is focused on the user's needs by understanding the data analysis task; (2) encode data using color and other visual characteristics; and (3) facilitate interaction by providing a direct manipulation user interface.

A number of well-known techniques exist for visualizing non-spatial datasets, such as, geometric projection, iconic display, hierarchical, graph-based, pixel-oriented and dynamic (or some combination thereof) [7, 8]. We decided an iconic display was most relevant to our goal of visualizing a memory chip's test data. Our visualizations were designed by first constructing an object to represent a single data element. Next, the objects are positioned to produce a static visualization of the memory chip test data. Glyphs are positioned based on scalar ranking attribute(s) within a traditional *2D grid* or along a linear *spiral* embedded in a plane.

3.1 Placement Algorithm

Glyphs representing the attribute values embedded in a dataset have to be positioned appropriately in order to create an information workspace for visual sense making. We decided to use two layout methods: a traditional 2D grid, and a spiral.

A two-dimensional ordering is imposed on the data elements through user-selected scalar attributes. We chose a 2D grid layout because it is intuitive and well-known placement algorithm.

A one-dimensional ordering is imposed on the data elements through a single user-selected scalar attribute, or "ranking" attribute. One way to map this ordering to a 2D spatial position is to use a 2D space-filling spiral. Our algorithm is based on a technique introduced by Carlis and Konstan to display data along an Archimedean spiral [9]. We have previously used 2D grid and spiral layouts to visualize storage controller performance data [10].

3.2 Data-Feature Mapping

When we design a visualization, properties of the dataset and the visual features used to represent its data elements must be carefully controlled to produce an effective result. Important characteristics that must be considered include [11]: (1) dimensionality (number of attributes in the dataset), (2) number of elements, (3) visual-feature salience (strengths and limitations that make it suitable for certain types of data attributes and analysis tasks), and (4) visual interference (different visual features can interact with one another, producing visual interference; this must be controlled or eliminated to guarantee effective exploration and analysis).

Perceptual knowledge of how the human visual system "sees" different properties of color and texture allow us to choose visual features that are highly salient, both in isolation and in combination [12–14]. We map visual features to individual data attributes in ways that draw a viewer's focus of attention to important areas in a visualization. Our glyphs support variation of spatial position, color and texture properties, including:

x*-position* and y*-position* or *linear radial position*, *hue*, *luminance*, *height*, *size*, and *orientation*. A glyph uses the attribute values of the data element it represents to select specific values of the visual features to display. After consulting with the domain experts we identify the attributes to include in a default data-feature mapping. The most important attributes should be mapped to the most salient features. The order of importance for the visual features we used is luminance, hue, and then various texture properties [12]. *ChipViz* allows the user to interact with the visualizations by translating, rotating, and zooming the environment. Users can change which visual features are mapped to each attribute using click-and-drag sliders. Finally, users can select individual data elements to display a pop-up balloon that describes the exact attribute values encoded by the element.

4 Visualization of Single Memory Chip Test Data

We selected different cases of test results taken from actual *DDR2/DDR3* products, and analyzed and interpreted the results with the help of our visualization tool, *ChipViz*. These cases represent four typical scenarios an engineer would encounter while analyzing memory chip test results. The first case shows the analysis of a lot with high fallout at *LT* and *low Vdd*. The second case describes the situation where high fallout in *Retention* occurs at *HT* and *high Vdd*. In the third case, we represent a more complicated scenario where high fallout occurs at *HT* due to multiple critical attributes (*tRCD* and *tRP*). Finally in the fourth case, we show the test results for a stable, high volume product in which the tests/processes are already highly optimized and results exceed the required pass rates.

4.1 Visualizing *Low Vdd at LT*

This dataset represents memory chip test results taken from a lot with high fallout. By visualizing the test results run at *LT*, it is evident that numerous tests have high failure rates and the overall yield is low. It is not immediately clear from the spreadsheet data what attributes are causing this high fallout, however.

By using *ChipViz* in Figure 1, we take advantage of visualizing multi-dimensional elements. The *Failure Rate* of a test is directly proportional to x-position and height, so high failure tests are sorted and can be easily viewed. y-position represents the *Test ID* number. In addition to this spatial filtering process, we visualize additional critical attributes among the high failing tests. The engineers requested to visualize *Vdd* as a primary attribute, and *tRAS* as a secondary attribute. *Vdd* is redundantly mapped to luminance, hue, and size (dark to light, red to blue, and small to large for lower to higher values, respectively). *tRAS* is mapped to orientation (more counterclockwise twist for larger values). By displaying the primary critical attribute *Vdd* with the most salient visual features (luminance, hue, and size), it is immediately evident that most of the high failure tests are dark, red, and large, showing an inverse relationship between *Failure Rate* and *Vdd*. Low *Vdd* is a common characteristic for most of the high failure tests. This is an important piece of information as it could point to design weakness for low voltage. The same information is visualized along a spiral in Figure 1b with distance

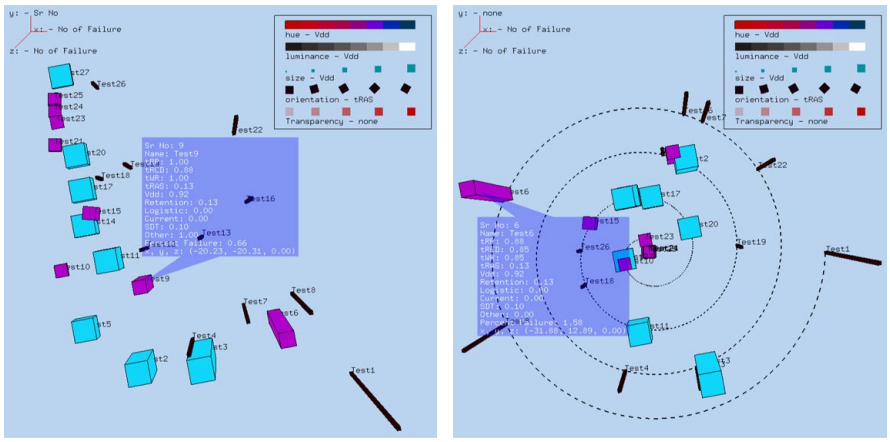

(a) (b)

Fig. 1. Visualizing *Low Vdd at LT*, *Vdd* → luminance, hue, size and *tRAS* → orientation: (a) *Failure Rate* → x-position, height and *Test ID* → y-position; (b) *Failure Rate* → radial position, height

from the center of the spiral proportional to *Failure Rate* (i.e., farther from the center for higher *Failure Rates*). We can easily conclude that as a data element moves away from the center of the spiral, its glyph becomes dark, red, and large, indicating the inverse relationship between *Failure Rate* and *Vdd*. Orientations varied randomly, suggesting no correspondence between *tRAS* and *Failure Rates*.

4.2 Visualizing *Retention at High Vdd and HT*

A second dataset is taken from a high volume memory chip tested at high temperatures. Retention at high temperatures is one critical attribute to test. For yield improvement, it is necessary to identify the top failing tests and identify their critical attributes. We can gain yield by trying to optimize these attributes. One of the most common problems for memory chips is *Retention* as the temperature increases. Again in Figure 2, we map the *Failure Rate* of the test to *x*-position and height, and y-position to *Test ID* in order to sort the data. We then map *Retention* to size. We can see that most of the high failure tests have large sizes, confirming a critical *Retention* component. By mapping *Vdd* to luminance and hue, we see most glyphs becoming bright and blue as we move along the x-axis, indicating that *Vdd* is directly proportional to *Failure Rate*. The high failure tests, apart from being *Retention* critical, also have high *Vdd*. This is confirmed by our spiral view. As we move away from the center of the spiral, the size of the glyphs increase, hue tends to blue, and luminance increases. As before, no patterns between *tRAS* (represented with orientation) and *Failure Rate* were visible.

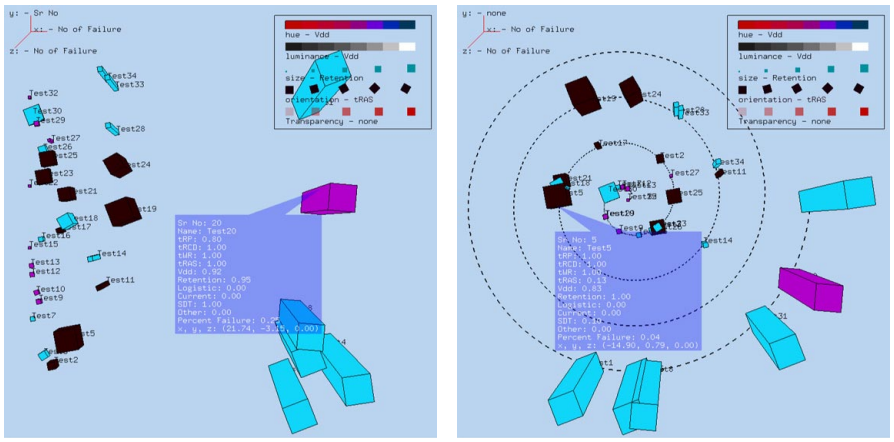

(a) (b)

Fig. 2. Visualizing *Retention* at *High Vdd and HT*, *Vdd* → luminance, hue, *Retention* → size, and *tRAS* → orientation: (a) *Failure Rate* → x-position, height and *Test ID* → y-position; (b) *Failure Rate* → radial position, height

4.3 Visualizing *tRCD, tRP at HT*

The third dataset represents a case with high fallout at high temperatures due to the row to column access delay (*tRCD*) and row precharge time (*tRP*) critical attributes. By looking at the spreadsheet of test results, it is difficult to decipher any useful information quickly, unless an experienced engineer remembers the critical attributes of every test. As before, *Failure Rate* is represented by *x*-position and height, and y-position represents *Test ID*. We map *tRP* to size and *tRCD* to orientation. From Figure 3, we see two trends: *tRP* decreases along the x-axis (smaller glyphs) and *tRCD* increases along x-axis (more counterclockwise twist). We can conclude that most of the high failure tests have two critical attributes, *tRP* and *tRCD*. In this dataset there was no correspondence between *Vdd* (represented by color) and *Failure Rate*. The same information can be gleaned from the spiral visualization.

4.4 Visualizing *Optimized Test Data*

The final dataset contains test results from a stable, high volume product. From the spreadsheet data, we can interpret that the *Failure Rate* for the tests are all low and of the same order. This usually happens for a product which is in high volume with optimum yields. There is a possibility that even though the *Failure Rate* for the tests is low, a particular attribute is contributing significantly toward the top failing tests. With *ChipViz*, we can try to interpret not only the individual *Failure Rate* of the tests but also the different attributes of each test by sorting the data and mapping each critical attribute to different visual features. In this dataset, from Figure 4 we do not see any correlation or trend among the different tests and their various attributes. This suggests

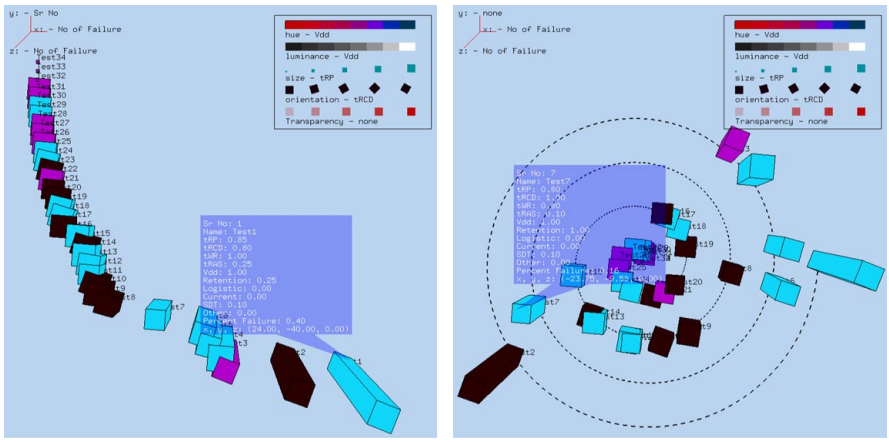

(a) (b)

Fig. 3. Visualizing *tRCD, tRP at HT, Vdd* → luminance, hue, *tRP* → size, and *tRCD* → orientation: (a) *Failure Rate* → x-position, height and *Test ID* → y-position; (b) *Failure Rate* → radial position, height

the dataset is for a product which is at a mature stage and for which the processes in the Front End (manufacturing sites) are stable and optimized.

5 Visualization of Multiple Memory Chips Test Data

One additional advantage of *ChipViz* is that we can visualize test data for more than one product. This is extremely helpful in analyzing test results for a product family or for products of a particular technology or from a particular manufacturing site. This allows engineers to identify design, technology or process issues on a much wider scale. Given this powerful capability, viewers can increase or decrease the resolution of the data analysis on the fly.

Figure 5 shows such a case for three types of memory chips visualized together. All three products belong to the same chip generation and technology, but have different physical sizes. We use *Failure Rates* to position each glyph, then define the chip being tested with *hue* (Figure 5a), *luminance* (Figure 5b), *size* (Figure 5c), or *orientation* (Figure 5d). The visualizations show that all three products have similar failure rates for most of the tests as expected. However, a small number of tests exhibit very different behavior across the different chips. That is, for some *Test ID* (i.e., rows in the visualization in Figure 5a), the failure rates for the three different chips are significantly different (e.g., there is no overlap between the three glyphs for *Test ID* 24). These tests are targeted for further analysis using individual chip visualizations from Section 4 to identify the common critical attributes that produce variable *Failure Rates*.

These examples illustrate the power and versatility of *ChipViz*. Our system can present complicated test results in a way that allows an engineer to decipher the results

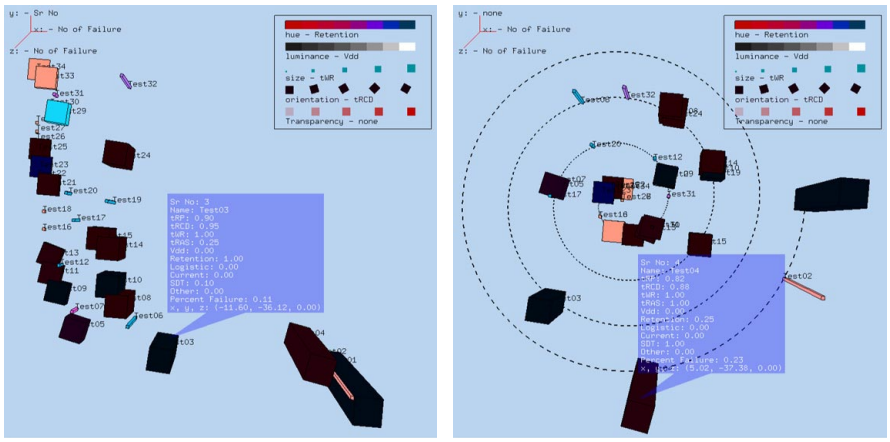

(a) (b)

Fig. 4. Visualizing optimized test data, *Retention* → hue, *Vdd* → luminance, *tWR* → size, and *tRCD* → orientation: (a) *Failure Rate* → x-position, height and *Test ID* → y-position; (b) *Failure Rate* → radial position, height

and draw conclusions efficiently. It also helps in highlighting information or relationships which are buried in the dataset.

6 Conclusions and Future Work

We have successfully applied perceptual visualization techniques to represent memory chip test data. This allows our engineering colleagues to gain more understanding of the relationships between various attributes measured during their testing process. Our results help the engineers rapidly analyze large amounts of test data and identify critical attributes that result in high failure rates. Based on anecdotal observations of real chip engineers, it took between one to two minutes to interpret a visualization. Our visualization methods are not necessarily restricted to memory chip test data, and may be useful for other datasets with appropriate ranking attributes. In the future, we would like to conduct validation studies to quantify our visualization design choices, and to measure the improvement our system provides over existing analysis techniques. We also plan to extend our techniques to analyze simulation results from VLSI circuit design and verification.

References

1. Oberle, H.D., Muhmenthaler, P.: Test pattern development and evaluation for drams with fault simulator ramsim. In: Proceedings of the IEEE International Test Conference on Test, Washington, DC, USA, IEEE Computer Society (1991) 548–555
2. Oh, S.C., Kim, J.H., Choi, H.J., Choi, S.D., Park, K.T., Park, J.W., Lee, W.J.: Automatic failure-analysis system for high-density dram. In: Proceedings of the IEEE International Test

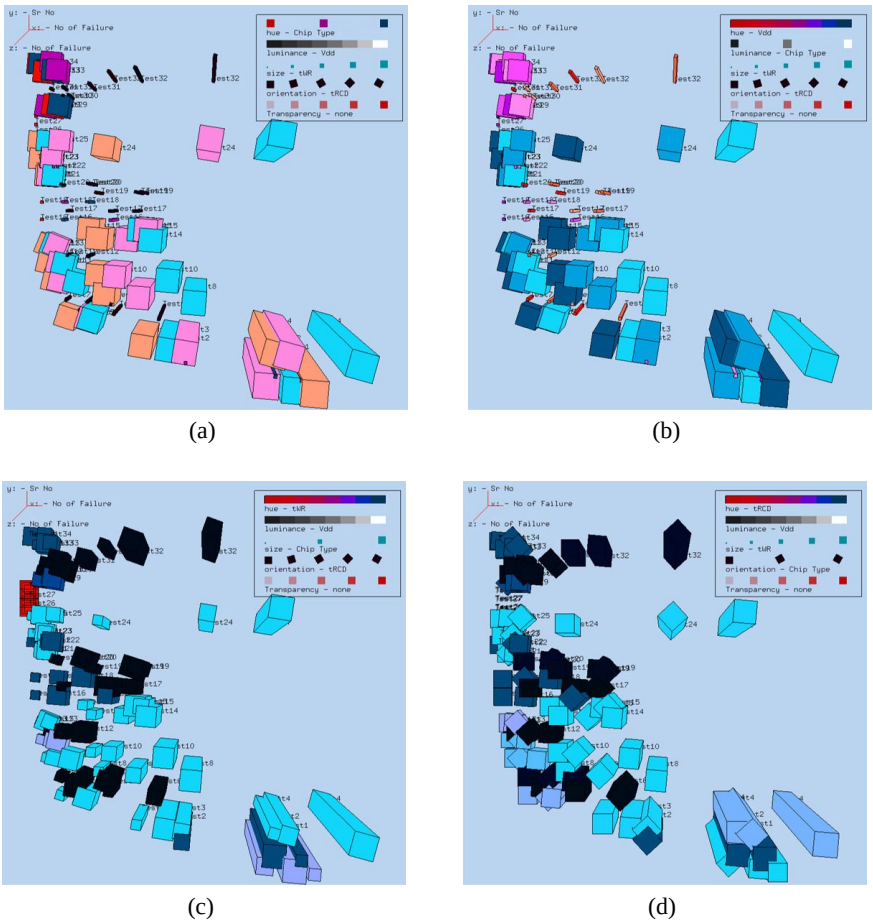

Fig. 5. Visualizing three types of memory chips, *Failure Rate* $\rightarrow x$-position, height and *Test ID* $\rightarrow y$-position, chip type mapped to: (a) hue; (b) luminance; (c) size; (d) orientation

Conference on TEST: The Next 25 Years, Washington, DC, USA, IEEE Computer Society (1994) 526–530

3. Fountain, T., Dietterich, T., Sudyka, B.: Mining ic test data to optimize vlsi testing. In: KDD '00: Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining, New York, NY, USA, ACM Press (2000) 18–25
4. Jones, J.A., Harrold, M.J., Stasko, J.: Visualization of test information to assist fault localization. In: ICSE '02: Proceedings of the 24th International Conference on Software Engineering, New York, NY, USA, ACM Press (2002) 467–477
5. van de Goor, A.J., de Neef, J.: Industrial evaluation of dram tests. In: DATE '99: Proceedings of the conference on Design, automation and test in Europe, New York, NY, USA, ACM Press (1999) 123
6. Eick, S.G.: Engineering perceptually effective visualizations for abstract data. In: Scientific Visualization, Overviews, Methodologies, and Techniques, Washington, DC, USA, IEEE

Computer Society (1997) 191–210
7. Keim, D.A.: Pixel-oriented database visualizations. SIGMOD Record (ACM Special Interest Group on Management of Data) **25** (1996) 35–39
8. Foltz, M., Davis, R.: Query by attention: Visually searchable information maps. In: Proceedings of Fifth International Conference on Information Visualisation, London, England (2001) 85–96
9. Carlis, J.V., Konstan, J.: Interactive visualization of serial periodic data. ACM Symposium on User Interface Software and Technology (1998) 29–38
10. Sawant, A.P., Vanninen, M., Healey, C.G.: PerfViz: A visualization tool for analyzing, exploring, and comparing storage controller performance data. In: Visualization and Data Analysis. Volume 6495, 07., San Jose, CA (2007) 1–11
11. Weigle, C., Emigh, W., Liu, G., Taylor, R., Enns, J.T., Healey, C.G.: Oriented texture slivers: A technique for local value estimation of multiple scalar fields. In: Proceedings Graphics Interface 2000, Montréal, Canada (2000) 163–170
12. Healey, C.G., Enns, J.T.: Large datasets at a glance: Combining textures and colors in scientific visualization. IEEE Transactions on Visualization and Computer Graphics **5** (1999) 145–167
13. Interrante, V.: Harnessing natural textures for multivariate visualization. **20** (2000) 6–11
14. Ware, C.: Information Visualization: Perception for Design, 2nd Edition. Morgan Kaufmann Publishers, Inc., San Francisco, California (2004)

Appendix: Definitions

Below are definitions of the attributes included in the memory chip test datasets:

1. *tRP* (Row Precharge time): the number of clock cycles taken between issuing a precharge command and an active command to the same bank
2. *tRCD* (Row Address to Column Address Delay): the number of clock cycles taken between issuing an active command and a read/write command to the same bank
3. *tWR* (Write Recovery time): the number of clock cycles taken between writing data and issuing a precharge command to the same bank, required to guarantee that all data in the write buffer can be safely written to the memory core
4. *tRAS* (Row Active time): the number of clock cycles taken between issuing an active command and a precharge command to the same bank
5. *Vdd*: power supply voltage
6. *Retention*: the maximum time a DRAM cell can store its programmed data
7. *Logistics*: a value defining whether there is a handling issue of chips in the manufacturing site
8. *Current*: current measured under different conditions on the chip
9. *LT*: low temperature (degree Celsius)
10. *HT*: high Temperature (degree Celsius)
11. *tRC* (Row Cycle time): the minimum time interval between successive active commands to the same bank, defined as *tRC* = *tRAS* + *tRP*