NPL National Physical Laboratory

Report

NPL Report
DEM-ES-015

Report to the National Measurement System Directorate, Department of Trade and Industry

From the Software Support for Metrology Programme

An Introduction to the Parallel Coordinate Technique Applied to Measurement Data

L Wright

NOT RESTRICTED

February 2007

dti National Measurement System

National Physical Laboratory | Hampton Road | Teddington | Middlesex | United Kingdom | TW11 0LW

Switchboard 020 8977 3222 | **NPL Helpline** 020 8943 6880 | **Fax** 020 8943 6458 | **www.npl.co.uk**

The DTI drives our ambition of ‘prosperity for all’ by working to create the best environment for business success in the UK. We help people and companies become more productive by promoting enterprise, innovation and creativity.

We champion UK business at home and abroad. We invest heavily in world-class science and technology. We protect the rights of working people and consumers. And we stand up for fair and open markets in the UK, Europe and the world.

The National Physical Laboratory is operated on behalf of the DTI by NPL Management Limited, a wholly owned subsidiary of Serco Group plc

An Introduction to the Parallel Coordinate Technique Applied to Measurement Data

L Wright
Mathematics and Scientific Computing Group

February 2007

ABSTRACT

Many measurement processes depend on a wide variety of factors such as sensor position, time, temperature, frequency, pressure, and so on. Hence the description of a measured value often includes specification of these extra factors, making the data set multi-variate. Similarly, the results of computer models are commonly large sets of multi-variate data. Multi-variate data sets are very common throughout metrology, and tools for visualising such data sets have a large range of potential applications.

Visualisation of multi-variate data is difficult: only three spatial dimensions are available for plotting, and usage of colour and symbols becomes difficult to interpret if too much information is presented simultaneously. Parallel coordinates offer an alternative approach to visualisation of such data sets.

This report aims to provide a brief introduction to the parallel coordinate technique, and to illustrate its application to real problems by using case study examples.

ISSN 1744–0475

National Physical Laboratory,
Hampton Road, Teddington, Middlesex, United Kingdom TW11 0LW

We gratefully acknowledge the financial support of the UK Department of Trade and Industry (National Measurement System Directorate)

Approved on behalf of the Managing Director, NPL
by Jonathan Williams, Knowledge Leader for the Electrical and Software team

Contents

1 Introduction

Many measurement processes depend on a wide variety of factors such as sensor position, time, temperature, frequency, pressure, and so on. Hence the description of a measured value often includes specification of these extra factors, making the data set multi-variate. Multi-variate data sets are common throughout metrology, and tools for visualising such data sets have a large range of potential applications.

As computing power becomes more readily available, mathematical models of physical processes can include complicated effects that had been neglected in earlier versions, leading to multiphysics models containing many variables. In addition, since models often run more quickly on faster computers, the time required to carry out Monte Carlo simulations and similar analyses has decreased, meaning that multivariate uncertainty analyses are now more tractable. Both of these developments result in large multi-dimensional data sets that often need to be interpreted visually.

Visualisation of multi-variate data is difficult: only three spatial dimensions are available for plotting, and usage of colour and symbols becomes difficult to interpret if too much information is presented simultaneously. Parallel coordinates offer an alternative approach to visualisation of such data sets.

Parallel coordinates were developed by Inselberg in the 1980s as a way of visualising multi-dimensional data sets. They have been successfully applied to a range of real problems, including process control and optimisation. They are particularly strong for identifying outliers, unexpected events, and the existence of multi-variable relationships.

This report aims to provide a brief introduction to the parallel coordinate technique, and to illustrate its application to real problems by using case study examples. Section 2 introduces the key ideas of the technique and suggests some possible applications. Section 3 describes how to interpret the plots, including techniques for getting the most out of each data set. Section 4 describes software implementations of the technique, including simple implementations within packages commonly used in metrology. Section 5 consists of case studies obtained from applying the technique to real problems. Section 6 suggests sources for further information and reading. Finally, an appendix illustrates of how some common features of two-dimensional measurement data appear in parallel coordinates.

It is important to note that the parallel coordinate technique is an interactive process, and so a report is not the ideal medium for the demonstration of the technique. Readers who are interested in further demonstrations are recommended to contact the author, or to investigate some of the software packages described in section 4. Additionally, please note that many of the figures in this report are easiest to interpret when viewed in colour, hence greyscale printed copies of the report may be more difficult to understand.

1.1 Vocabulary

The nature of the technique, and its use of multi-segmented lines to represent points in n-dimensional space, makes definition of some simple terms very important. Throughout this report, unless stated otherwise:

- A *point* means a point in n-dimensional space.
- A *line segment* means a part of a parallel coordinate representation of a point linking two axes.
- A *line* means a line in n-dimensional space.

1.2 Acknowledgements

This guide was produced under the Software Support for Metrology (SS*f*M) programme, which is managed on behalf of the DTI by the National Physical Laboratory. For more information on the SS*f*M programme, contact the programme manager, Mr Bernard Chorley (phone +44 20 8943 7050, email ssfm@npl.co.uk, National Physical Laboratory, Hampton Road, Teddington, Middlesex, UK TW11 0LW), or visit the website at http://www.npl.co.uk/ssfm.

The author would also like to thank Stephen Robinson and Gary Hayman, both of the Acoustics Team in the Division of Quality of Life at NPL for the use of their data in the first data set, the suppliers of the industrial data set presented in the second case study (who have asked to remain anonymous), and Robin Brooks of Curvaceous Software for his comments and suggestions for improvements.

2 The Parallel Coordinate Technique

2.1 What are parallel coordinates?

Suppose that we wish to explore a data set $\{\mathbf{x}_i, i = 1, 2, \dots, N\}$ where each of the $\mathbf{x}_i$ consists of n quantities (either inputs or outputs), so that $\mathbf{x}_i = \{x_i^j, j = 1, 2, \dots, n\}$. If n is 2, we can plot the two quantities against one another in an X-Y plot, and each of the $\mathbf{x}_i$ will be a point in two-dimensional space. If n is 3, it may be possible to visualise the data set in three-dimensional space, and if n is higher than 3 it may be possible to use colour or some other property of the point markers to visualise the other quantities. However, there is a limit to how many quantities can be visualised on a single plot without the plot becoming difficult to interpret. Even three-dimensional plots are difficult to understand if the data points do not lie on a regular smooth surface.

One alternative is to plot subsets of the data as two-dimensional plots. The observer then has to try and work out how these separate plots interact and what they mean, which is not straightforward. The number of quantities that can be visualised simultaneously is still limited. Additionally, if n is large, a large number of plots will be required to examine every combination of quantities.

In X-Y plots and three-dimensional visualisations, the axes along which quantities are plotted are perpendicular to one another. If we want to use perpendicular axes, we can only ever visualise a maximum of three quantities using spatial dimensions. If we want to visualise all of the quantities simultaneously using spatial dimensions, we must use parallel axes and the parallel coordinate technique.

The parallel coordinate technique is a way of displaying all of the multi-dimensional data set $\{\mathbf{x}_i, i = 1, 2, \dots, N\}$ at once. The axes representing the quantities are drawn as parallel rather than perpendicular. In the simplest form of the technique, each axis is normalised so that the plotted quantity is actually

$$\hat{x}_i^j = \frac{x_i^j - x_{\text{min}}^j}{x_{\text{max}}^j - x_{\text{min}}^j} \tag{1}$$

where $x_{\text{min}}^j = \min_{i=1,2,\dots,N}\{x_i^j\}$ and $x_{\text{max}}^j = \max_{i=1,2,\dots,N}\{x_i^j\}$. This normalisation means that the scales on different axes cannot be compared directly, so for instance a 3 on one axis may not be in the same position as a 3 on another axis. The data points are then represented by drawing straight line segments connecting the axes at the points that make up $\mathbf{x}_i$. Each set of connected straight line segments represents a data point.

As an example, suppose we want to visualise the following five six-dimensional points:

$$
\begin{aligned}
\mathbf{x}_1 &= \{80, & 97, && 86, && 13, && 94, && 86\},\\
\mathbf{x}_2 &= \{3, & 48, && 57, && 52, && 78, && 56\},\\
\mathbf{x}_3 &= \{43, & 67, && 26, && 85, && 97, && 94\},\\
\mathbf{x}_4 &= \{89, & 74, && 79, && 13, && 72, && 87\},\\
\mathbf{x}_5 &= \{32, & 66, && 94, && 76, && 25, && 7\}.
\end{aligned}
$$

First the points would be normalised using the formula given above, which would give the new points as

$$
\begin{aligned}
\hat{\mathbf{x}}_1 &= \{0.90, & 1.00, && 0.88, && 0.00, && 0.96, && 0.91\},\\
\hat{\mathbf{x}}_2 &= \{0.00, & 0.00, && 0.46, && 0.54, && 0.74, && 0.56\},\\
\hat{\mathbf{x}}_3 &= \{0.47, & 0.39, && 0.00, && 1.00, && 1.00, && 1.00\},\\
\hat{\mathbf{x}}_4 &= \{1.00, & 0.53, && 0.78, && 0.00, && 0.65, && 0.92\},\\
\hat{\mathbf{x}}_5 &= \{0.34, & 0.37, && 1.00, && 0.88, && 0.00, && 0.00\},
\end{aligned}
$$

to two decimal places. Most parallel coordinate software packages will carry out the normalisation automatically, but if a non-specialist package is being used then the user will need to normalise the data. Finally the data points are represented as line segments linking the parallel axes. The visualisation of this data set is shown in Figure 1. Reading the axis representing Q1, the first quantity, from top to bottom, the line segments represent $\mathbf{x}_4$, $\mathbf{x}_1$, $\mathbf{x}_3$, $\mathbf{x}_5$, and $\mathbf{x}_2$.

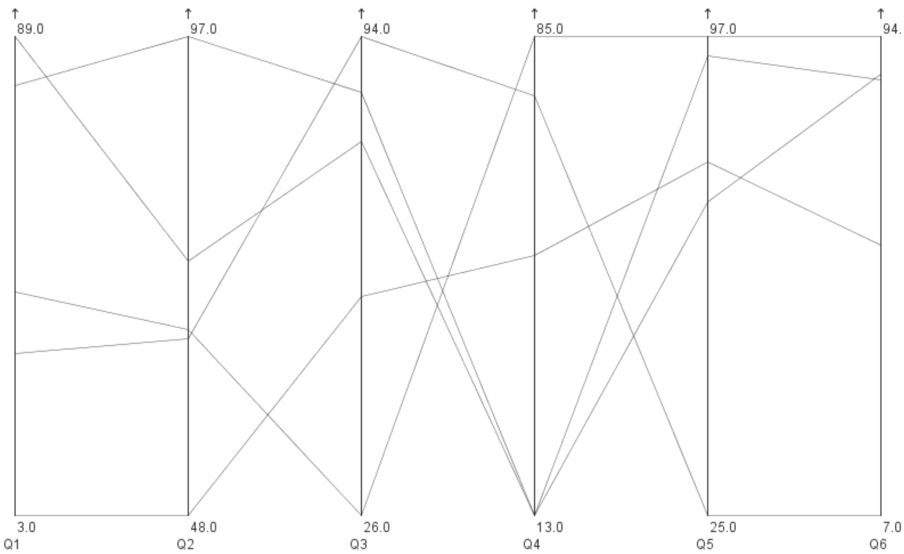

Figure 1: A simple example of a parallel coordinate plot.

The line segments representing each point are easy to identify in this example. For instance, the line segments that link the top end of the final three axes must represent $\mathbf{x_3}$ because the final three normalised coordinates of $\mathbf{x_3}$ are 1.00. Identification of individual points becomes more difficult when more points are plotted, but some software packages have tools to overcome this difficulty.

The two most useful features of the parallel coordinate technique are interactivity and the ability to identify multi-dimensional relationships qualitatively. These features will be discussed further in section 3.

2.2 Applications of parallel coordinates

Figure 1 shows a very simple example that only includes five points, and the points were generated at random so no trends are observable within the data. Any conclusions that can be drawn from this plot could probably be drawn by examining the numerical data by eye. However, as N, the number of data points, increases, drawing conclusions from the raw data becomes a lot more difficult. For data sets consisting of a large number of points and a large number of dimensions, the parallel coordinate technique offers more of an advantage than it does for small data sets of low dimension.

In particular, the parallel coordinate technique can be used to identify trends and relationships in broad terms so that the quantities involved in the trends can be examined in more detail using other techniques.

Parallel coordinates are also extremely useful for large-dimensional data sets. In particular, if a problem has a large number of input parameters and a smaller number of output parameters, and it is not clear how these two are related, parallel coordinates make visualisation of the relationship between the inputs and the outputs more simple.

A common example of this type of problem is process control. Manufacturing processes often have a small number of 'output' quantities, such as scrap rate, product quality measures, and total output, but most multi-step processes have several controllable quantities at each step (e.g. temperature, material flow rate, etc.), giving a large number of input quantities in total. If the aims are to reduce scrap rate, produce high quality product, and increase total output, it is rarely clear what values need to be assigned to the various controllable quantities to achieve these aims. Visualisation using parallel coordinates can help to identify ranges of values for the inputs, and it can indicate which of the inputs need to be tightly controlled.

This type of problem has some equivalents in metrology. In general, the metrologist wishes to reduce uncertainties, and has a number of input quantities that can be controlled during the measurement process, such as temperature, humidity, location of sensors, and so on. A parallel coordinate visualisation of these input quantities and the corresponding uncertainties could identify the experimental conditions under which the uncertainties are smallest.

As an example, consider a system with five input quantites: two temperatures T_1 and T_2, two pressures p_1 and p_2, and a flow rate Q, and a single output quantity U. Suppose the input and output quantities are linked by the equation

$$U = \min(1, \max(Q\frac{p_1 - p_2}{T_1 - T_2}, -1)) \tag{2}$$

so that U always lies between -1 and 1, and that all the input variables are randomly chosen and normalised to lie between 0 and 1. Plotting the output quantity against each of the input quantities results in an unhelpful graph, an example of which is shown in figure 2. Similarly, plotting U as a surface function of two variables is not useful, as illustrated by figure 3.

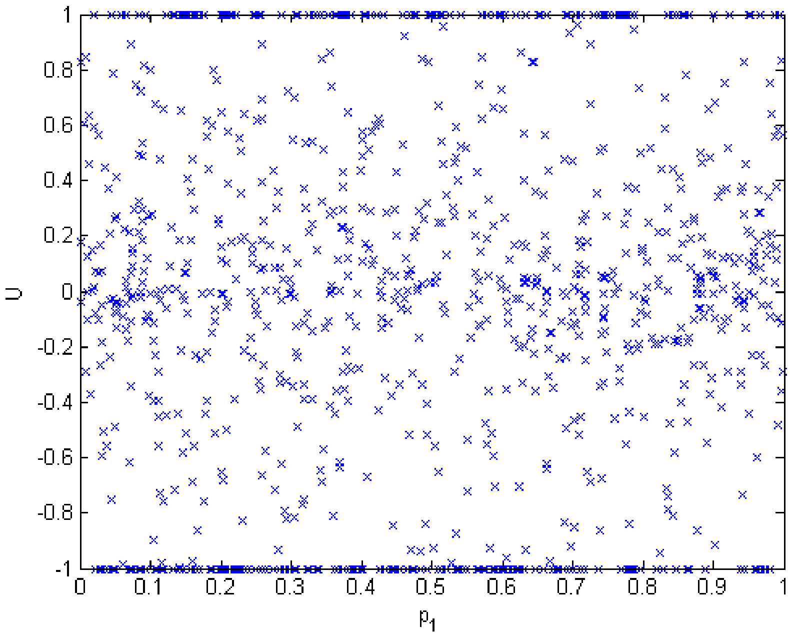

Figure 2: Output variable U plotted as a function of the input variable p_1.

However, if these points are plotted as a parallel coordinate plot the data become easier to explore and hence easier to understand. The three plots shown in figures 4, 5, and 6 show points selected to illustrate some of the dependencies of the output on the input. Each highlighted subset contains about 10% of the points.

Note that all of the input variables in figures 4, 5, and 6 are plotted on axes ranging from 0 to 1. The figures show that the flattest line segments between T_1 and T_2 tend to be associated with values of U of 1 and -1, the flattest line segments between p_1 and p_2 tend to be associated with values of U close to zero, and the smallest values of Q also tend to produce values of U close to zero. These observations agree with the equation given above, and could be used to motivate further plots. For instance, the observations about the flattest line segments could lead to plots using new variables $dp = p_1 - p_2$ and $dT = T_1 - T_2$.

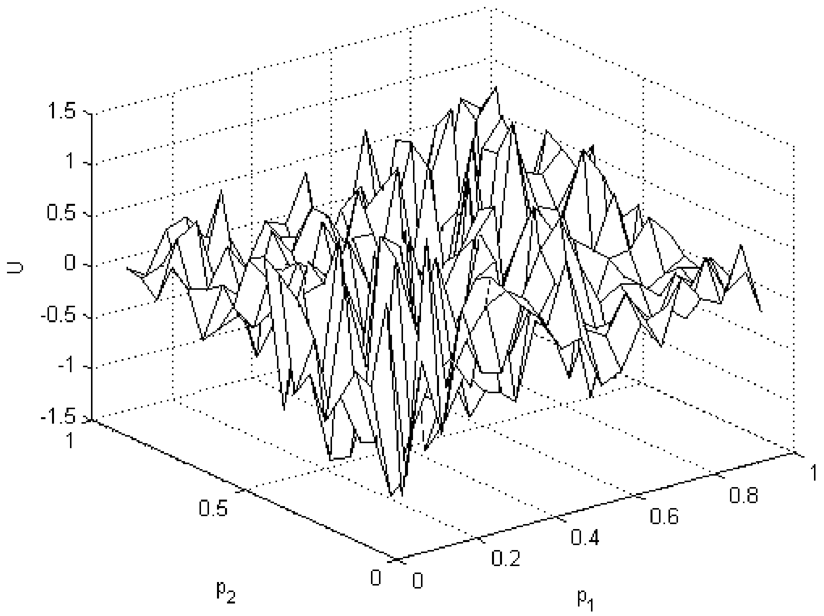

Figure 3: Output variable U plotted as a function of the input variables p_1 and p_2.

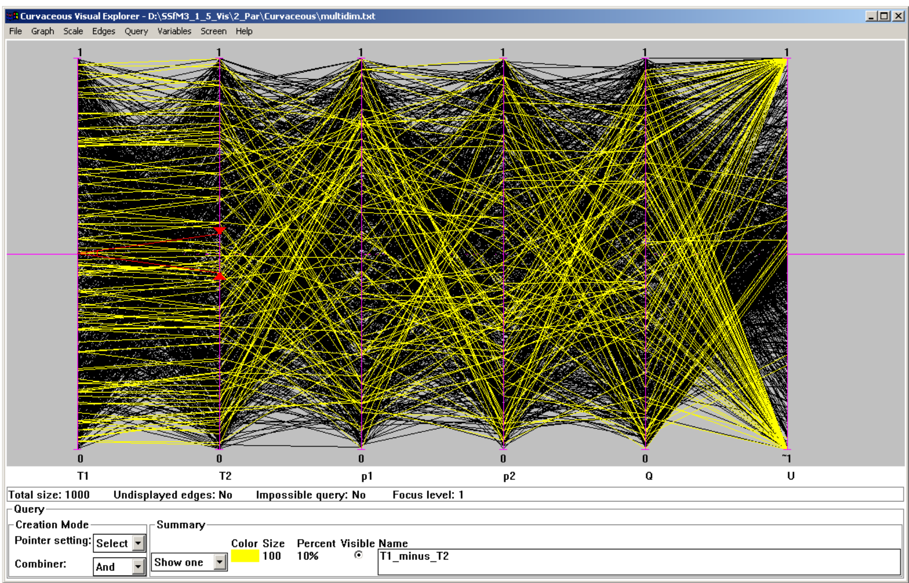

Figure 4: Flattest line segments linking T_1 and T_2 are highlighted.

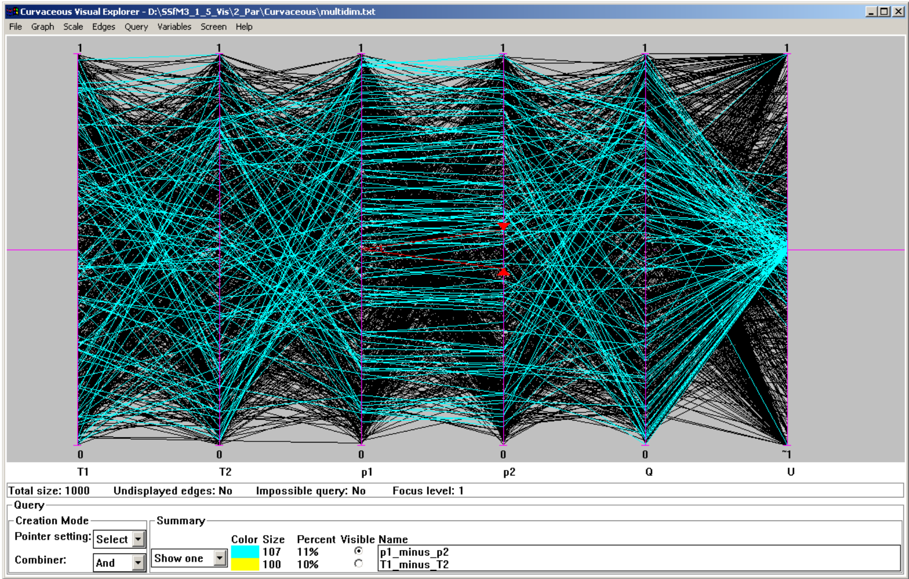

Figure 5: Flattest line segments linking p_1 and p_2 are highlighted.

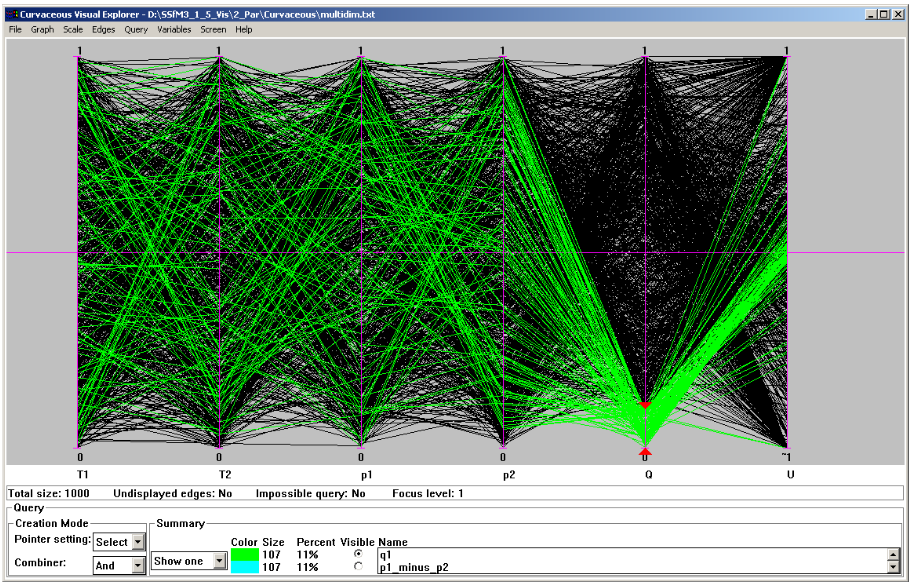

Figure 6: Smallest values of Q are highlighted.

2.3 Limitations of the technique

There are several limitations to the parallel coordinate technique. The main limitation is on a fundamental property of the data. If each of the variables only takes a small, finite number of values, the plots can be very difficult to interpret. As an example, suppose that the data set contains six input variables (x_1 to x_6) and one output variable (y), and that each of the input variables can take one of three values (-1, 0 and 1, say). Then if the data set is generated by testing every possible set of input quantities, 729 data points will be generated. If $y = \sum_{n=1}^{6} 0.1 n x_n$, then figure 7 shows a typical plot. The discretised nature of the data makes it hard to draw any conclusions from this simplistic form of the parallel coordinate plot.

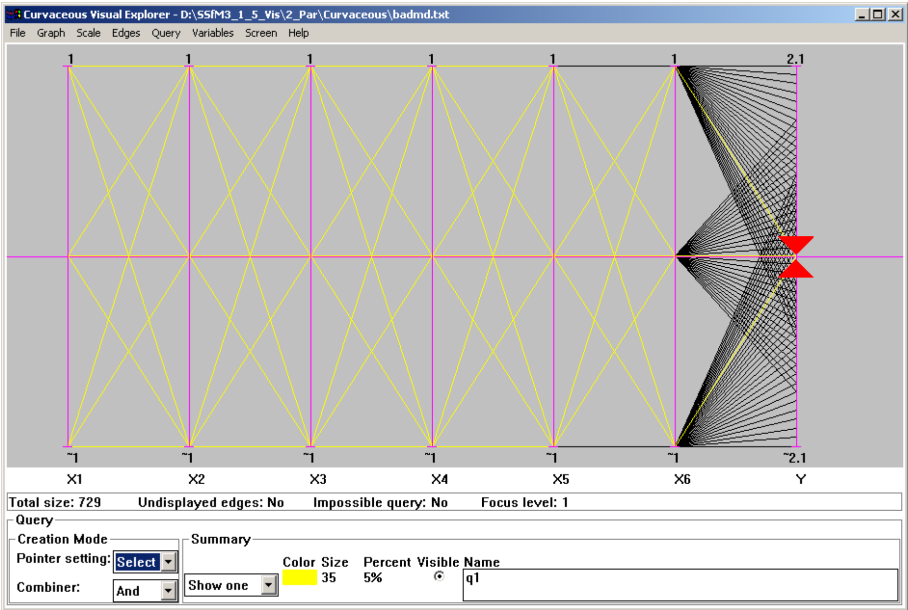

Figure 7: An example of a data set that is unsuitable for the parallel coordinate technique.

This is an over-simplified example, but this type of data set is commonly encountered in Designed Experiment activities and in applications of the Response Surface Methodology. Various tools exist for visualisation of response surfaces, including some based on displaying every possible scatter plot formed from choices of pairs of variables and one based on the parallel coordinate technique.

Similarly, data sets consisting of a small number of points may not be suitable for the technique since each variable will only take a small number of values. It is always difficult to identify trends and relationships in a high-dimensional data set containing a small number of points, no matter what method is used for visualisation.

There are other limitations on the data sets the technique should be applied to. For two- and (to a lesser extent) three-dimensional data sets, the technique offers little or no advantage over traditional techniques and hence should probably not be used. Another (minor) limitation is that direct comparison between axes is lost: a 3 on one axis is not necessarily a 3 on another. This limitation is a matter of learning a new way of interpreting the plots, and can be remedied if necessary by setting all axes to the same scale.

Another barrier to the use of the technique is that it effectively requires a new way of interpreting visualisations: the plots are not intuitive and familiar like a two-dimensional line. However, the new interpretative skills are simple and easy to pick up so this problem is a minor one.

Another barrier concerns choice of software: the technique is extremely limited if the software used to generate the plots does not allow interaction. See section 4 for more discussion on this point.

3 Interpreting parallel coordinate plots

This section explains how to interpret parallel coordinate plots. At first, the plots can seem confusing as they contain a lot of data in a form that is unfamiliar. Here we present some tips on how to get the most out of the plots. As an extension, appendix 1 illustrates some features of parallel coordinate plots that relate to common features of measurement data.

3.1 General tips

The tips in this section can be easily applied if a parallel coordinate visualisation tool is being used. If another tool is being used to implement the technique, such as Excel or Matlab, most of the tips can be implemented by appropriate sorting, transforming or scaling of the data. For instance, assuming that the procedure outlined in section 2.1 has been followed so that the data have been scaled to lie between 0 and 1, an axis can be reversed by plotting $1 - x$ instead of x. Similar transformations can be carried out to produce the other effects described below.

However, it should be noted that one of the key features of the parallel coordinate technique is its interactivity: it allows data to be explored as a dynamic process rather than by examining static pictures. Use of a generic software package generally loses this interactvity, and so wherever possible a parallel coordinate visualisation tool should be used for implementation.

3.1.1 Pairwise comparisons

The parallel coordinate technique makes pairwise comparisons simple without having to plot every possibly pair of variables of interest as an X-Y plot. To compare two variables, put their axes next to one another in the plot. If a parallel coordinate tool is being used, there is usually a tool for permuting axes. If some other tool is being used, it may be necessary to reorder the columns of the matrix that contains the data and plot the data again.

The number of pairwise comparisons that can be made between n axes is $\frac{1}{2}n(n-1)$. As n increases, it becomes more important to choose carefully the comparisons that are made, whilst not ignoring any potentially useful effects. It is possible to design an optimal set of permutations of the axes that includes every possible pairwise comparison in the minimum number of plots. This set is called a "grand tour", and some parallel coordinate tools calculate an appropriate set of permutations automatically so that users can cycle through the permutations to examine their data thoroughly. For example, suppose that measurements are made of three quantities A, B, and C, and that the measurements are dependent on time t, temperature T, and length L. Then all of the possible pairwise comparisons can be made if the axes in the plots are ordered as shown in Figure 8.

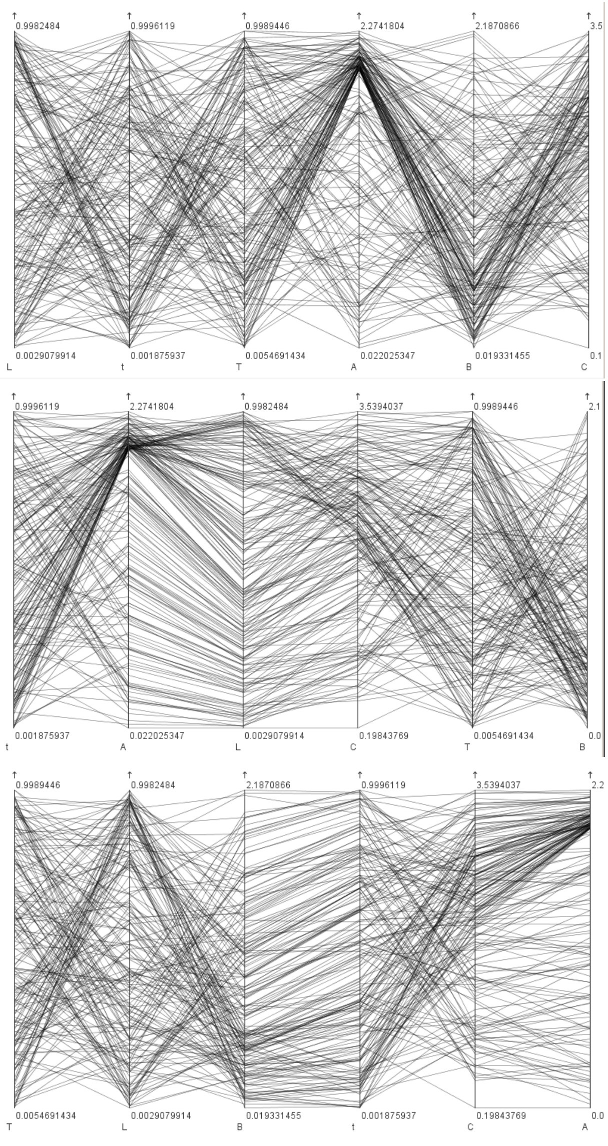

Figure 8: All possible pairwise comparisons of six variables are displayed in these three plots. Axes are ordered $\{L, t, T, A, B, C\}$, $\{t, A, L, C, T, B\}$, and $\{T, L, B, t, C, A\}$.

It should be noted that the grand tour concept is usually not particularly useful. In general, the data set being visualised has a known set of inputs and a known set of outputs, and the aim of the visualisation is to explore how the inputs affect the outputs. This knowledge means that the only pairwise comparisons worth making are between inputs and outputs (and possibly outputs and outputs to look for correlation).

It is interesting to note that a permutation of the axes is equivalent to a rotation in n-dimensional space.

3.1.2 Using subsets of the data

Some parallel coordinates tools make it possible to highlight subsections of the data so that groups of data with common features can be traced. This can make identification of key features of the data considerably easier. If the implementation of the parallel coordinate technique being used does not allow for this highlighting, it may be worthwhile to produce plots of subsets of the data in order to investigate further. This technique is particularly useful for data sets consisting of many points, as will be shown in the examples (see Section 5).

Highlighting tools are very useful for identifying relationships that involve more than one variable, a key feature of the parallel coordinate technique. The relationships that can be identified are only ever qualitative, but once a relationship has been identified, further analysis of the data can take place to produce quantitative relationships.

As an example, consider the plot of the data set used in Figure 8 shown in Figure 9(a). There is a particularly dense set of lines on the A axis (the fourth axis from the left). Figure 9(b) shows the same data with the subset of interest highlighted, and the plot makes it clear that the lines in the dense set all have L values in the upper half of the axis.

More complicated relationships can be identified by choosing more complicated subsets. For example, the plot in Figure 10 was generated by selecting the lines that pass through the lowest 30% of the A, B, and C axes. It shows that if A, B, and C are to be small then L and t must be small, but that T is not really important. This information could be useful if low values of A, B, and C were important but it was not known which of the inputs should be controlled.

3.1.3 Manipulating the axes

As is explained in section 2, most parallel coordinate plots initially scale the axes linearly between the minimum value in the data set at the bottom and the maximum value in the data set at the top. This is not always the easiest form for identification of relationships between data sets. Inverse relationships are easier to identify if one of the axes is reversed, and if a subset of the data is of interest then focussing the axis on one particular range of values is

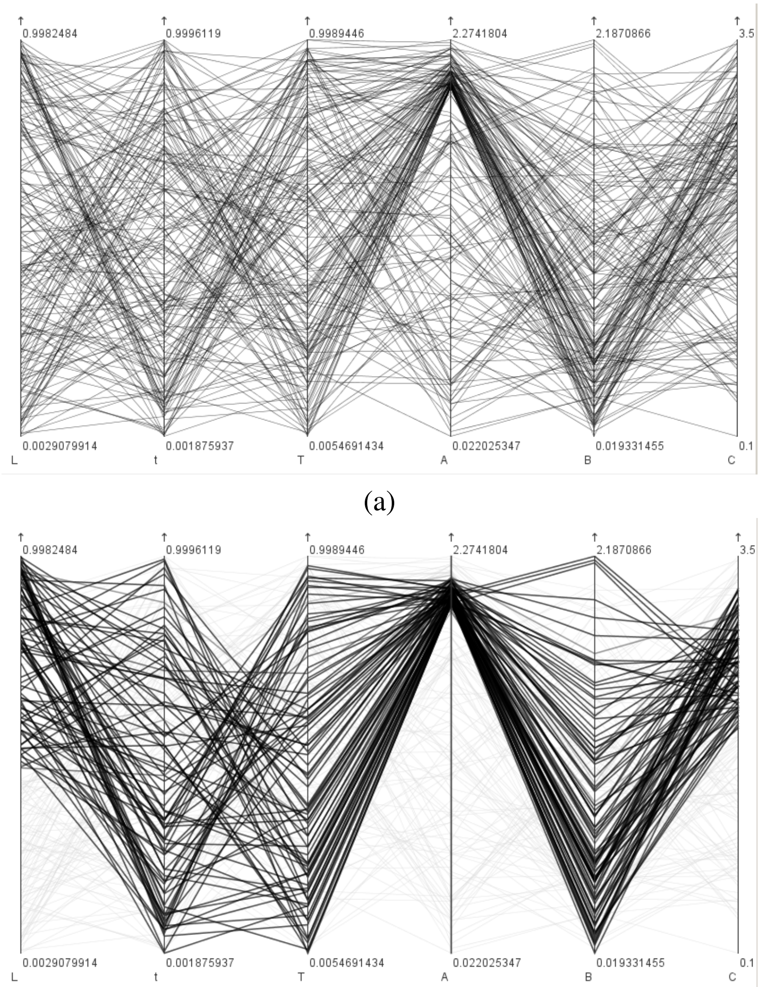

(a)

(b)

Figure 9: The full data set, (a), and a highlighted subset, (b).

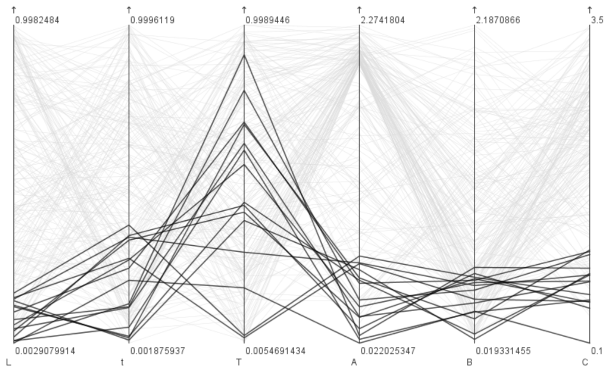

Figure 10: Data set with line segments that pass through the lowest 30% of the A, B, and C axes highlighted.

likely to be a useful technique. If several variables represent the same quantity measured in different positions it can be useful to ensure that the axes representing those variables are all to the same scale so that intercomparison is possible. In most parallel coordinate tools the axes can be rescaled to zoom in on an area of particular interest or to enable simple comparisons to be made.

For example, the A axis in Figure 9(b) can be manipulated so that the highlighted set of points can be examined more closely. The resulting plot is shown in Figure 11. Whilst this plot is still not easy to interpret, it is more clear than the previous representation. A careful highlighting of lines shows that the highest values of A have high values of t and T, and the values of A closer to 2.0 have at least one of t and T small. Figure 12 shows two examples: Figure 12(a) shows that for similar values of L and T, A is proportional to t, and Figure 12(b) shows that for similar values of L and t, A is proportional to T. Further examples of these trends can be found.

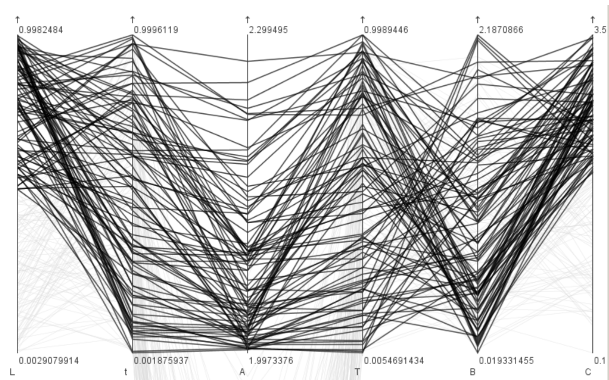

Figure 11: Data that is highlighted in Figure 9(b) plotted with the A axis rescaled.

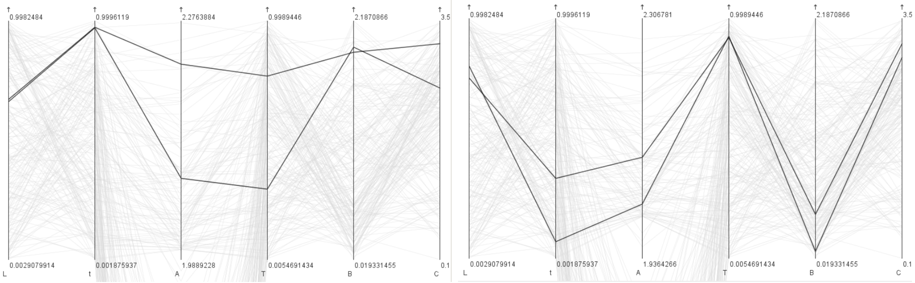

(a) (b)

Figure 12: Two pairs of points from Figure 11 highlighted, showing the dependency of A on T and t.

3.1.4 Multi-dimensional relationships

The key strength of the parallel coordinate technique is its ability to visualise multi-dimensional relationships. However, utilisation of this strength requires the user to change the way they think about visualisations. Multi-dimensional relationships are not shown in a way that is an extension of the common two- and three-dimensional line and surface plots.

The parallel coordinate technique enables users to explore multi-dimensional data sets interactively. The technique allows the user to choose a set of conditions on a subset of the variables (usually the output variables) and asks which values of the other (input) variables generate that subset. This can enable the user to see how to control input variables in order to generate output variables in the required range. Conversely, examining the output variables corresponding to chosen subsets of the input variables can give insight into the likely dependencies.

As an example, consider the data set discussed in section 2.2, where

$$U = \min(1, \max(Q\frac{\Delta p}{\Delta T}, -1)) \tag{3}$$

where Q lies between 0 and 1, and $\Delta p = p_1 - p_2$ and $\Delta T = T_1 - T_2$ lie between -1 and 1. Figure 13 shows the data set with all the negative values of U highlighted. Note that the ΔT axis has been reversed. The main point of interest is that none of the line segments between the Δp and ΔT axes cross the (imaginary) line segment connecting the zeroes of both axes. This feature indicates that all of the data points with negative values of U had values of Δp and ΔT of opposite signs. A similar conclusion can be drawn linking the positive half of the U axis and points which have the same sign for Δp and ΔT. Further plots can be created, by restricting the slope of the highlighted line segments, that show that the size of U is approximately proportional to the size of Q.

Examples of how common two-and three-dimensional closed shapes look are given in section 3.3, and their extensions to higher dimensions is discussed in brief. For interest,

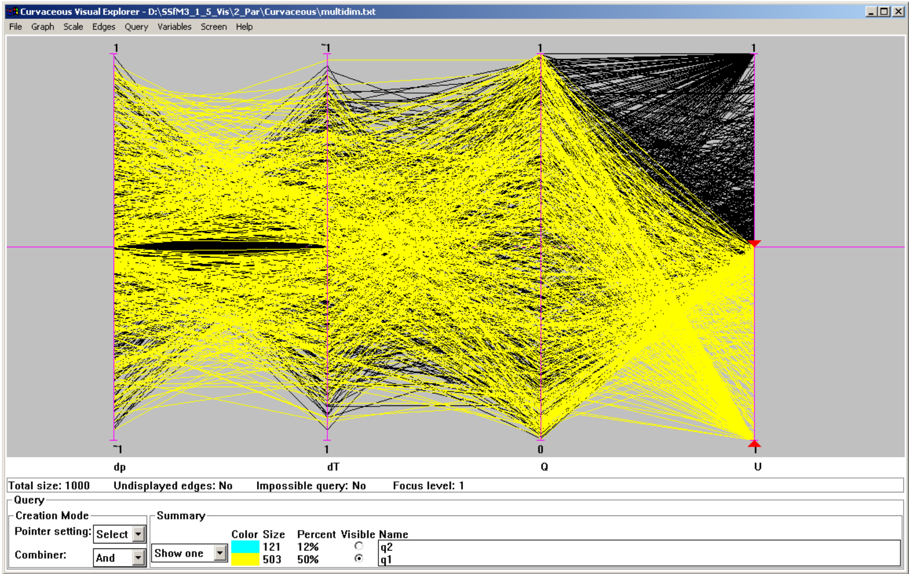

Figure 13: Points with negative values of U are highlighted.

figure 14 shows a sphere in 6-dimensional space. The curves that form at the upper and lower limits between axes are typical, and the “ideal” limit that would be generated with an infinite number of curves can be shown to be of this shape.

A large number of papers exist that prove various similar geometric results linking shapes in Cartesian coordinates with their equivalents in parallel coordinates. The simplest result is that a straight line in parallel coordinates is equivalent to a point in Cartesian coordinates, and vice versa. This is obvious for a two-dimensional line but if parallel coordinate line segments are allowed to extend beyond their axes it can be shown for higher numbers of dimensions (see section 6 for references). Whilst these results are useful when using parallel coordinates for automated control systems, they are generally of less assistance when trying to interpret a plot, so they will not be discussed here.

3.2 Summary

- Explore the multi-dimensional relationships interactively by selecting subsets of the data dynamically.
- Comparing every pair of data sets that may be of interest may be worth while.
- Reverse the ordering of one of the axes to check for inverse relationships that may not be clear as a direct visualisation.
- Rescale the axes to focus in on areas of particular interest or to highlight any relationships that are hidden with the normal scaling.

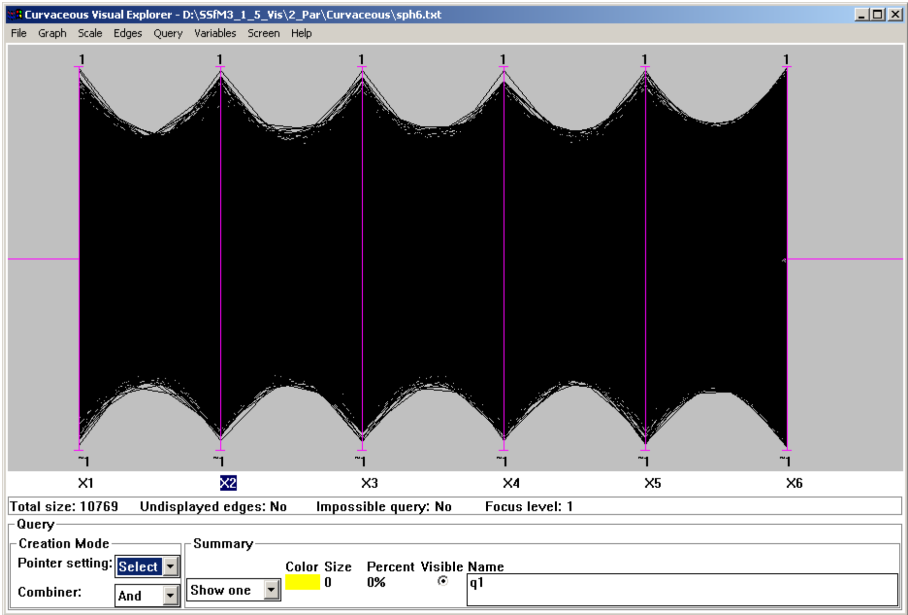

Figure 14: Parallel coordinate visualisation of a 6-dimensional sphere.

- Examine subsets of the data (by highlighting, recolouring, or replotting) to get more insight into detailed relationships.

3.3 Common two-dimensional shapes

It is useful to consider how some closed-form two-dimensional shapes look when plotted as parallel coordinates, since the two-dimensional forms lead to a clearer understanding of the n-dimensional equivalents.

The following figures were generated by randomly creating 5000 points in the region $-1 \leq x_1, x_2, \leq 1$ and assigning them with a value of 1 if they lay within the shape of interest and a value of 0 otherwise. The shapes of interest were:

- a square with $-0.5 \leq x_1, x_2, \leq 0.5$,
- a rectangle with $-0.5 \leq x_1 \leq 0.5, -0.75 \leq x_1 \leq 0.75$
- a circle of radius 0.5 with $x_1^2 + x_2^2 \leq 0.25$,
- an ellipse of half-axes 0.75 and 0.5 with major and minor axes lying along the x_1 and x_2 axes so that $\frac{x_1^2}{0.75^2} + \frac{x_2^2}{0.5^2} \leq 1$,
- the same ellipse rotated through 45° so that $\frac{(x_1+x_2)^2}{0.75^2} + \frac{(x_1-x_2)^2}{0.5^2} \leq 2$,

- a torus of inner radius 0.25 and outer radius 0.5, so that $0.125 \leq x_1^2 + x_2^2 \leq 0.25$.

Note that all figures are shown with the x_1 and x_2 axes scaled between -1 and 1. The x_1 and x_2 axes are the first two axes, and the line segments of interest throughout are the line segments connecting these two axes.

Figure 15 shows the points that lie within the square. Note that all of the line segments lie with a region bounded by straight lines. Due to the scaling, the bounding lines are horizontal. The rectangular region is shown in figure 16, and again the region is bounded by straight lines. If the x_2 axis was scaled to have limits of -1.5 and 1.5, these bounding lines would be horizontal because 0.5 on the x_1 axis would be the same level as 0.75 on the x_2 axis.

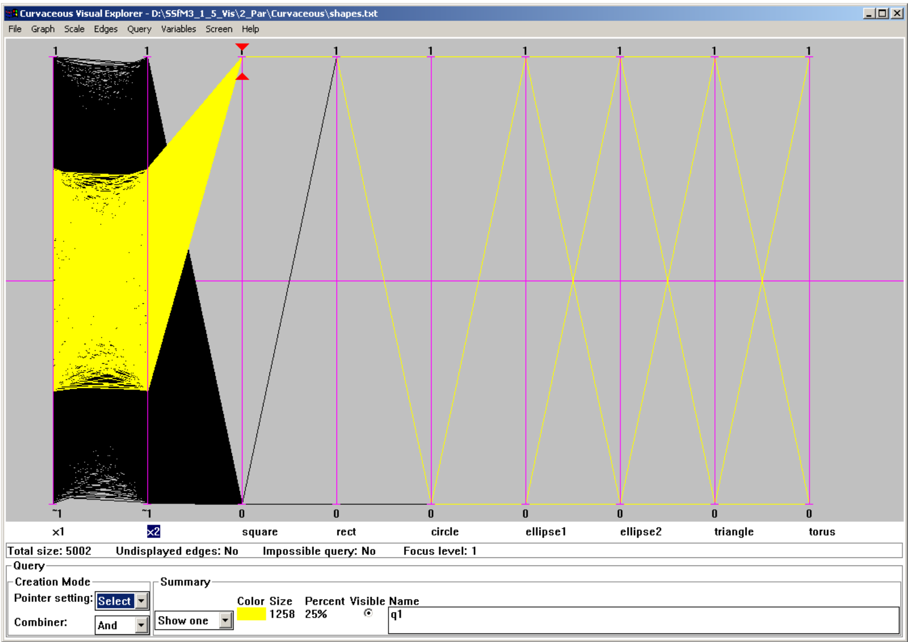

Figure 15: Parallel coordinate plot highlighting points within a square region.

Figure 17 shows the points lying within the circle. The curved bounding lines have the characteristic shape shown in figure 14. Figure 18 shows the points lying within the unrotated ellipse, and the bounding lines are similarly curved. An expression can be derived for this bounding line in terms of the limits of the scales of the axes and the half-axes of the ellipse.

Figure 19 shows the points lying within the rotated ellipse. The curved bounding lines here resemble those of the circle, but it can be shown that they are not quite identical. This illustrates one of the difficulties of interpreting parallel coordinate plots. Another difficulty is illustrated in figure 20, which shows the torus points. At first glance, the highlighted points look exactly the same as those in figure 17, but they are not the same. Figure 21

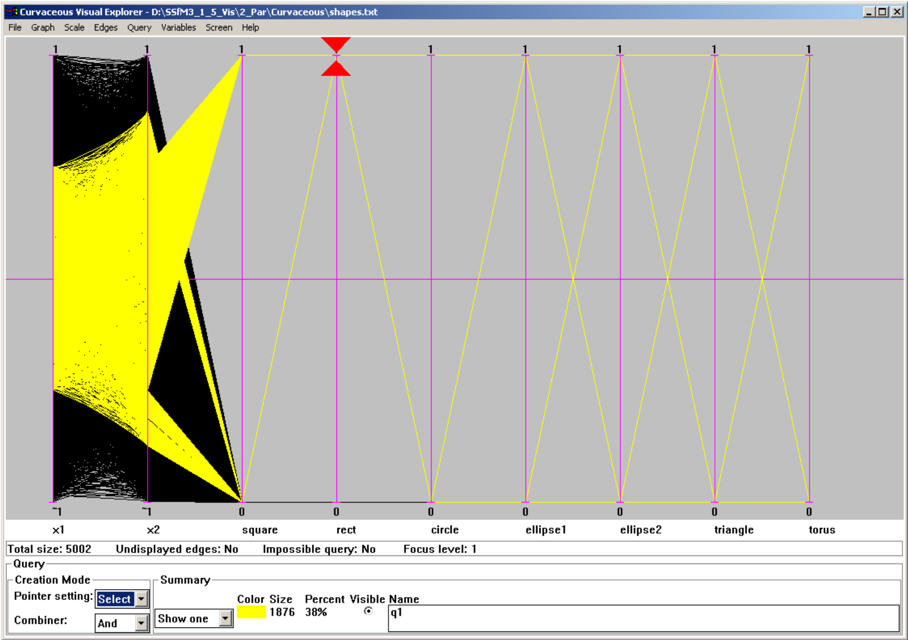

Figure 16: Parallel coordinate plot highlighting points within a rectangular region.

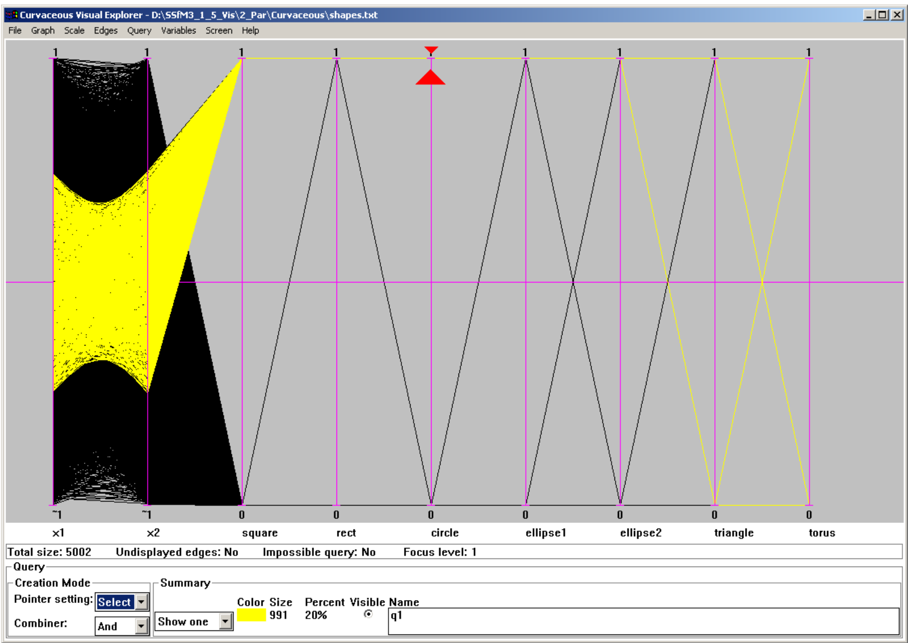

Figure 17: Parallel coordinate plot highlighting points within a circular region.

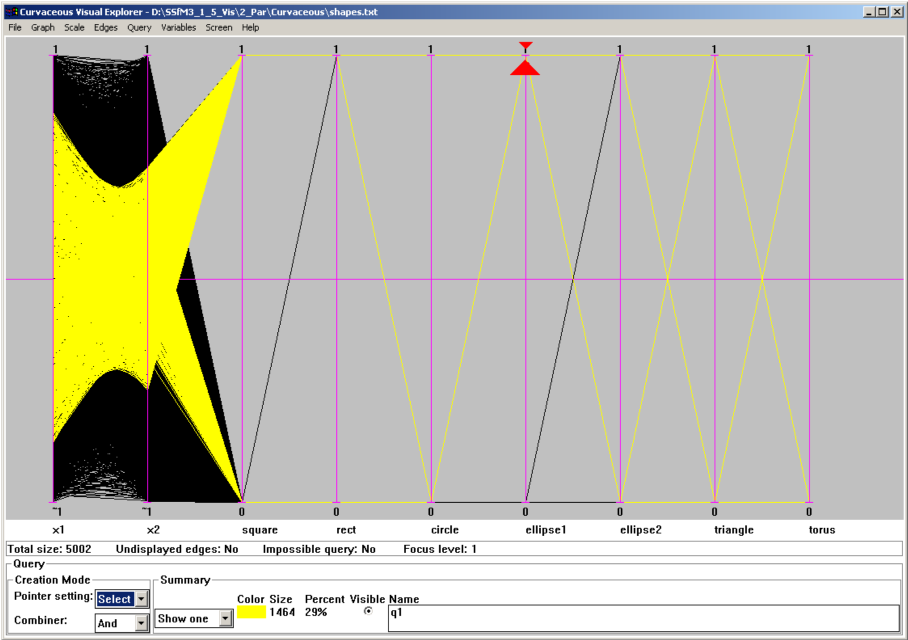

Figure 18: Parallel coordinate plot highlighting points within an elliptical region.

highlights the points that are in the torus in yellow and those that are in the circle but not the torus in blue. The parallel coordinate technique can disguise holes in the data set.

3.4 Extensions

The main extension of the parallel coordinate technique is its application to control systems. Parallel coordinates make it possible to define control limits on variables and define the limits of variable interaction so that automated warning and control systems can be set up to act whenever the process approaches the known limits of operation. The parallel coordinates control method is better han applying limits to each variable independently because it allows for interactions, and hence produces tighter limits, without requiring a formula linking the inputs and outputs.

As an example, suppose that there are two input variables x_1 and x_2, and one output variable y, and that $y = 1$ if $x_1^2 + x_2^2 \leq 0.25$ and $y = 0$ otherwise, as shown in figure 22. The aim of the controls is to ensure that $y = 1$ at all times. The simplest approach, assuming that nothing is known about how the three variables are related, is to apply independent limits to x_1 and x_2. The black line in figure 22 represents the maximal limits that must be applied to x_1 and x_2, but it encloses some areas where $y = 0$, so the limit is not tight enough. The white line is the widest limits that can be applied if the condition $y = 1$ is to be maintained, but it excludes some areas where $y = 1$, which is inefficient.

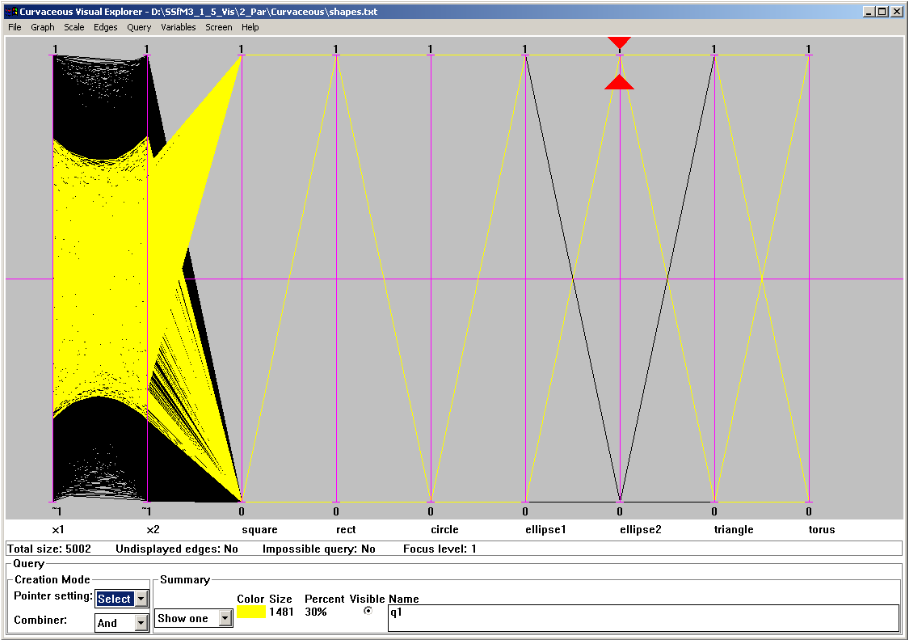

Figure 19: Parallel coordinate plot highlighting points within a rotated elliptical region.

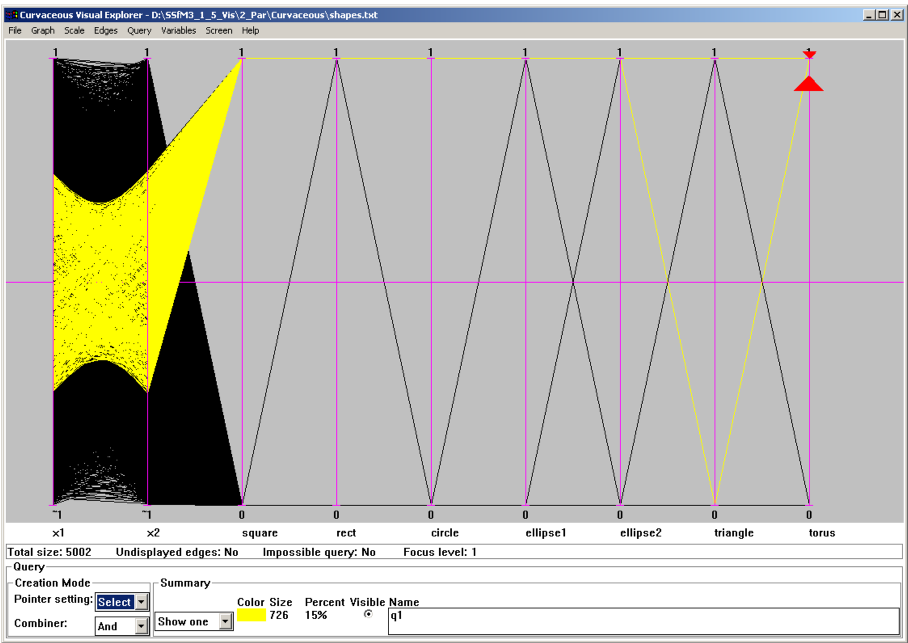

Figure 20: Parallel coordinate plot highlighting points within an toroidal region.

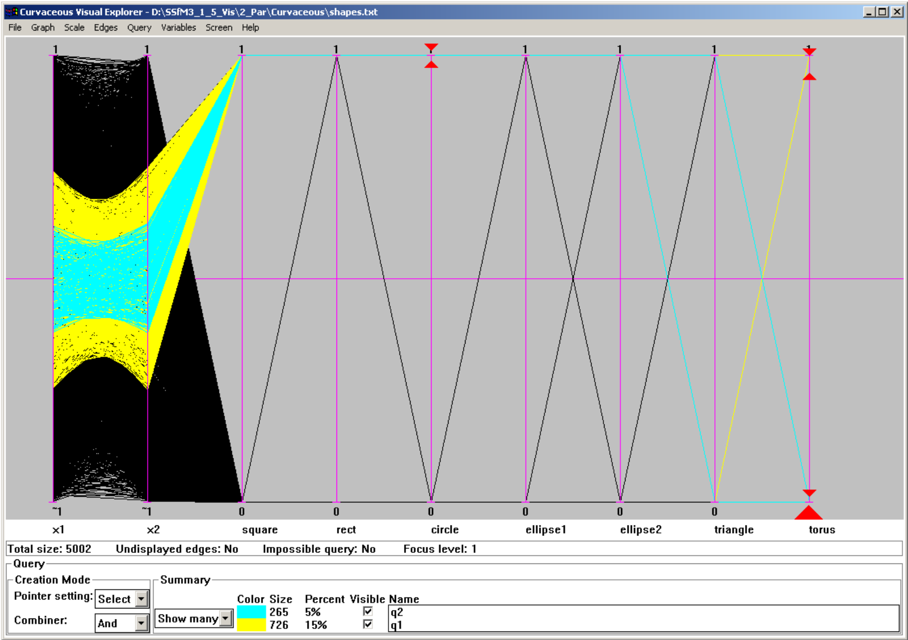

Figure 21: Parallel coordinate plot highlighting points within a toroid and within its inner circle separately.

Figure 23 shows the same problem as a parallel coordinates plot. The highlighted points are those with $y = 1$. Note that the highlighted line segments between the x_1 and x_2 axes have the same characteristic shape as the circle in section 3.3. If restrictions are placed on the input variables such that the input variables have to remain within the yellow area, the output will always be $y = 1$, and the yellow area covers the whole of the region where $y = 1$ rather than a smaller square within it.

This is a trivial example because it is a two-dimensional example and a formula is known that links the inputs and the output. In multiple dimensions, the approach of applying limits on each variable seperately restricts the “acceptable” volume hugely. Whilst the maximally inscribed square as shown in figure 22 is has about 64% of the area of the circle shown, the maximally inscribed hypercube in six dimensions has less than 6% of the volume of the equivalent hypersphere. This makes applying limits in each direction independently very wasteful in multiple dimensions. Trying to fit multi-dimensional hypersurfaces to experimental data is difficult, particularly if the data points are sparse and irregularly spaced, and so a process control methodology that avoids the need for such fitting but does not over-restrict the variables is immensely valuable.

This methodology has a potential application to metrology. It is common to express the multi-dimensional equivalent of confidence intervals as multi-dimensional ellipsoids. In two dimensions it is (reasonably) straightforward to draw conclusions about the correlation and sensitivity of the variables by looking at the half-axes and the angles between the

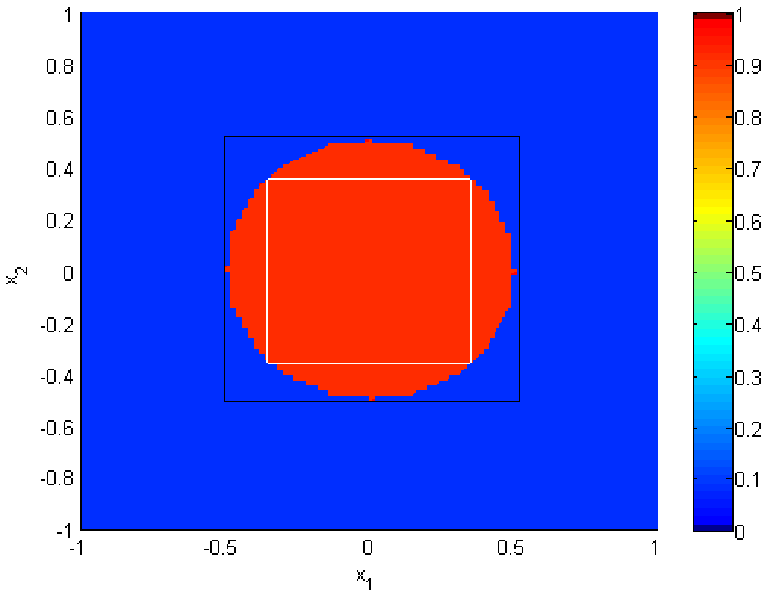

Figure 22: Plot showing variation of y, and two possible methods of imposing limits on x.

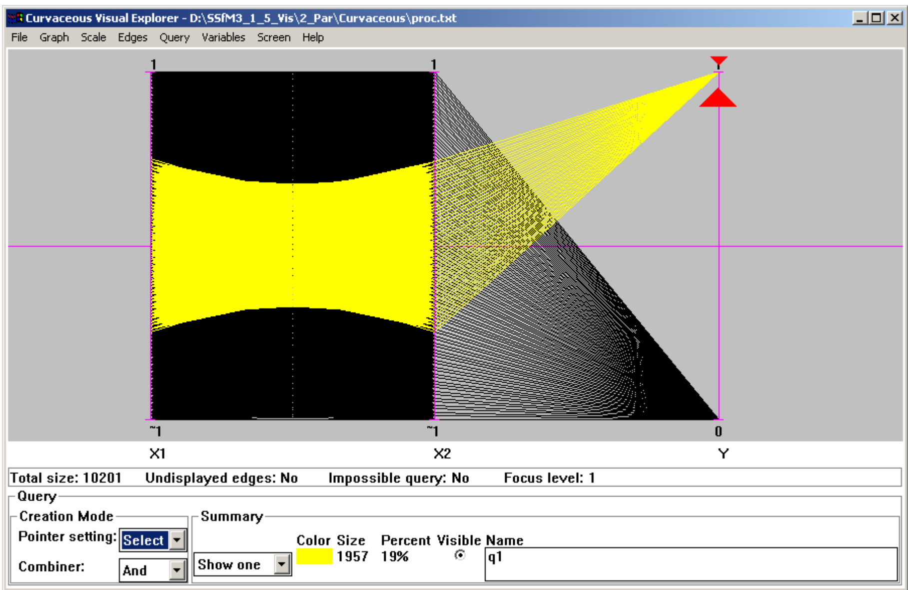

Figure 23: Parallel coordinate plot of the data, with lines showing $y = 1$ highlighted.

ellipse's axes and the variable axes. However, in higher dimensions this becomes more difficult to do. It is possible that the parallel coordinate technique could provide a good way of visualising these multi-dimensional ellipsoids.

4 Software Implementations

This section describes some of the software tools that are available for visualisation of data using the parallel coordinate technique. In order to evaluate them, it is necessary to identify which fetaures are important in a good software implementation of the technique.

As was mentioned in section 2, the two most useful features of the parallel coordinate technique are interactivity and the ability to identify multi-dimensional relationships qualitatively. Any good software implementation of the technique should make these two features easily accessible. Multi-dimensional relationships are easiest to identify interactively, by considering how one or more variables change as the values of the other variables are controlled. Even if pictures illustrating data features are viewed statically, they will be generated interactively. Hence ease of interaction is an important point to consider when judging parallel coordinate software.

The various highlighting, reordering, and scaling tools mentioned in section 3 are extremely useful features of software implementations. As the technique and the associated tools are relatively new to many potential users, clear documentation and preferably worked examples and a sample data set are important. As the software tool is implementing a visualisation technique, the ability to create image files of the results of the visualisation process is handy.

4.1 Software packages

Two main software packages have been used to generate the figures in this report: a freeware package called Parvis and a commercial package called C:Suite Visual Explorer (CVE). Please see section 6 of this report for information on how to obtain all packages mentioned in this section.

Parvis has been used to generate the plots in sections 2 and 3. Parvis is a Java package that uses mouse interactions to perform the various operations described in section 3. It accepts data in a “standard table file” format, a tab-separated format that is easy to prepare from spreadsheets or other applications that generate tabulated data.

Parvis offers the full range of manipulation tools mentioned in section 3, including the ability to combine “brushes” (the tools for highlighting subsets of the data) to carry out complex Boolean operations. The package also offers a “correlation brush”, which highlights lines showing positive or negative correlations between a user-defined pair of variables. This can be useful for identifying relationships between axes that are not adjacent. Parvis can also overlay histograms on each of the axes, which can be a useful tool for understanding how the variables are distributed when there are too many lines overlapping for interpretation to be straightforward.

The main drawback with the tool is that the documentation is somewhat patchy (probably because the package is freeware). Whilst the tool is easy to use, it is not always obvious what some of the options do (in particular the working of the correlation brush is largely unexplained).

CVE has been used to generate the figures in section 5. CVE is a commercially available software package that uses the parallel coordinate technique. The package was originally developed for applications in process control (the developers were originally process engineers), and has been used in many industrial applications. The tool can be treated as a stand-alone package (accepting data in comma delimited, space delimited or tab delimited formats) or it can be started from within Excel, so that data in a spreadsheet can be read directly into CVE. The tool can save the plots that it creates as graphics files (bitmaps or portable network graphics files), and can save the current state of a visualisation in its own format so that users can return to partly-developed visualisations.

The tool includes all of the manipulation tools described in section 3, and has a number of extensions of each of these tools. The rescaling operations available include a group rescaling tool, so that similar variables can be plotted to the same scales and thus compared directly, a zero-centred scaling tool, and an interactive option to alter each scale to user-selected limits. Similarly, the permutations tools for reordering variables include a full set of permutations that comprise a grand tour (see section 3.1 for a definition) as well as the option to create user-defined permutations.

As well as these standard tools, CVE includes several other special tools. The tools for highlighting subsets of the data set can choose ranges of a given value, or lines of a given slope, or several other options. Once a subset has been highlighted, the non-highlighted data can be removed in order to focus in on the points of interest. Whilst data is highlighted the tool indicates how many lines have been chosen, how many lines there are in the full set, and what the former is as a percentage of the latter. Users can also define new variables from simple mathematical functions of existing variables.

Additionally there are several more complicated variables that can be created by analysing the data and looking for clusters, calculating the derivative of a series, or sorting the data in the same order as some combination of variables. The tool also offers X-Y scatter plots of the data, including any highlighting, so that users can examine pairwise interactions more closely.

In order to help the user explore the data set more effectively, the software includes a "Box query" tool. This tool is particularly useful in process control. Suppose that the data being visualised consists of input variables $x_1, x_2, ..., x_m$ and output variables $y_1, y_2, ..., y_n$. If ranges of each of the y have been selected, then each of the x will also lie within some range (possibly the full range of the variable, possibly not). For example, in figure 10 the ranges of A, B, and C have been treated as the output variables and as a result, the ranges of L and t have been restricted. These ranges of the input variables define a box in m-dimensional space. The box query tool highlights every point on the parallel coordinate plot that lies within that box. These points show what would happen to the output variables if the input variables were controlled to lie within the box, and so the box query can help to identify the benefits of changing the control regime of a process.

The tool is supported by extensive documentation and the developers offer training courses. Once the initial training has been undertaken, it is an intuitive tool to use, and offers the full range of tools required to create and interpret parallel coordinate plots.

In addition to these packages, the authors are aware of several other visualisation packages with parallel coordinate capabilities.

Xmdv (available from http://davis.wpi.edu/ xmdv/) is, according to its website, ”a public-domain software package for the interactive visual exploration of multivariate data sets”. It features four main types of plot: scatter plots, star glyphs, parallel coordinate plots, and dimensional stacking. The authors have found the tool significantly less easy to use than Parvis and CVE. In particular, the highlighting and reordering tools recommended in section 3 did not seem to be available and the data highlighting tools seemed less intuitive than those in other packages.

XGobi is a package designed for visualising multi-dimensional data whose main tools use linked line and scatter plots that allow for data highlighting to enable a user to look at different two-dimensional slices of the same data set simultaneously. The tool also includes some parallel coordinate capability, but the capability is only a small part of the tool and has not been extensively developed.

Orange Canvas is a wide-ranging set of freeware Python-based tools for visualisation that includes a parallel coordinates capability. The plots produced using this tool do not seem to be interactive, which limits their effectiveness. Crystal Vision is similarly wide-ranging in terms of the range of tools offered but limited in terms of its parallel coordinate capabilities.

Another freeware package, CASSATT, is available for Mac and Windows PC, but the authors were unable to find a downloadable Windows version and so have not used the tool.

A simple internet search has yielded a number of other commercial packages that support parallel coordinate plots, including Stata and S, and some developed solely for parallel coordinate visualisations (including Parallax, written by the originator of the technique, Professor Alfred Inselberg).

4.2 Implementations in other common packages

As was explained in section 2.1, parallel coordinate plots are usually based on a simple linear rescaling of the data such that each quantity is normalised to lie between 0 and 1 for maximum visual resolution, although some packages allow for logarithmic rescalings. Hence it is possible to plot all of the quantities relative to a single y axis in any software package that can manipulate data and plot line graphs, such as Excel or Matlab.

Matlab features a parallel coordinates tool in its Statistics Toolbox. The tool is not interactive, and does not carry out normalisation in the way described in section 2, probably because quantities are being treated like random variables. If requested, the quantities are normalised such that they have mean 0 and standard deviation 1 before visualisation. The most interesting feature of the tool is its ability to plot percentiles of the data by linking the n^{th} value on each axis (although this feature is misleading for data visualisation since the line thus produced will not necessarily represent a data point).

The use of such packages has a number of disadvantages. In particular, the interactivity and data exploration capability is lost. Not all of the tools offered by most parallel coordinate packages are available, which can make the interpretation of data sets containing a large number of points difficult. Additionally, since the points are all plotted relative to a single y axis, the range of each plotted quantity is not shown as it is in most parallel coordinate plots.

It may be possible to use Matlab or Excel to carry out a small feasibility test to check whether the parallel coordinate technique is useful for a given data set. However, it is likely that the lack of interactivity will lead to subtleties of the data being missed and the technique being wrongly rejected.

The most beneficial use of such implementations of the parallel coordinate technique is to help new users understand the mathematics underlying the technique, and to help them to understand the move from perpendicular axes to parallel ones. In particular, visualisation of familiar mathematical functions and surfaces can aid understanding and increase familiarity before a more interactive tool is used to visualise real data.

4.3 Recommendations

A good software implementation of the parallel coordinate technique should make interactive data exploration simple. Ideally the implementation should include the manipulation tools mentioned in section 3 as well as good documentation and the ability to create image files.

In the author's experience, freeware packages are available that can give a good initial insight into multi-dimensional data sets. Most packages have the required interactivity and many have the manipulation tools, but not all are easy to use. The lack of documentation and limited functionality means that commerical tools are more suitable for more complicated investigations.

The commercial tool that the author has used has the full range of interactive functionality that a good implementation needs, and can visualise several different complicated queries simultaneously, making comparisons simple. It also has extensive documentation and support, and creates image files easily. Commercial packages are suitable for visualising complicated features of large multi-dimensional data sets.

Implementations of the parallel coordinate technique in packages commonly used for visualisation, such as Matlab or Excel, are not interactive and so are of limited use for real data. However, such implementations can provide a useful tool for understanding the mathematics underpinning the technique and for introducing new users of the technique to the idea of parallel axes.

5 Case Studies

5.1 Case Study 1: Hydrophone calibration data from an outdoor facility

5.1.1 The data set

NPL's Wraysbury Acoustic Calibration Laboratory is an open-water test and calibration facility. The facility consists of a calibration raft in the centre of a 1 km by 2 km fresh water reservoir with an average depth of 22 metres. On the raft are three separate, fully instrumented measurement stations capable of determining the full range of acoustic parameters that characterise the performance of acoustic transducers or complete sonar systems. The facility includes cranes and handling equipment capable of deploying large sonar arrays. The service offered by Wraysbury is UKAS-accredited for the frequency range 1 kHz to 350 kHz, and free-field measurements are also possible at frequencies below 1 kHz.

As a part of its measurement capabilities, the facility has a set of reference transducers that cover the frequency range 1 kHz to 350 kHz. The range is covered by four sets of transducers, with nominal operating frequencies 8 kHz (range of 1 – 20 kHz), 50 kHz (15 – 85 kHz), 150 kHz (70 – 220 kHz) and 300 kHz (200 – 350 kHz). The first three frequency ranges have four different transducers in each set, and the final frequency range has a set of three transducers.

Each time a set of measurements is made, the conductance and sensitivity of a transducer are measured for a set of frequencies in the operating range of the transducer. The same set of frequencies is used for all transducers in a set. As well as measuring the transducer response, the date is noted and the temperature of the water is measured. Not all of the transducers in a set are measured on the same date. The intervals between measurement dates vary from 35 days to 330 days, and data sets are available back to 1995.

5.1.2 Nature and organisation of the data

There are three types of parallel coordinate plot in this case study. The first type, an example of which is figure 31, was created by plotting the data gathered from a single set of transducers (i.e. all devices with the same frequency range) on a single set of axes. The axes chosen were device number, date (treated as an integer using the Excel system of days elapsed since 1st January 1901), month, water temperature, frequency, capacitance and sensitivity. This choice of axes provides a good overview of the data sets.

The second type of plot, of which figure 24 is an example, has been created by taking three key frequencies from the operating range (the maximum, the minimum, and the nominal operating frequency), and plotting device number, date, month, water temperature, and sensitivity and conductance at each of the three frequencies on separate axes (10 axes in total). This provides a good way of visualising possible relationships between variables, since plotting all the frequencies simultaneously can mask the subtleties.

The third type of plot, developed to compare devices with one another, uses axes of frequency, date, month, water temperature, and conductivity and sensitivity of each of the devices separately. If a device was not measured at a given date and frequency, the conductance and sensitivity of the device are assumed to be the minimum of the existing values on that axis minus 5% of the range of the axis (e.g. if an axis runs from 0 to 20, any missing data would be plotted at -1). See figure 26 for an example.

5.1.3 Questions to be answered by the visualisations

The metrologists who gather these data sets have a number of questions that they would like to investigate.

1. Are there any relationships between measured values and temperature?
2. Are there any relationships between measured values and date? In particular, do any of the devices drift over time?
3. Is there any frequency-dependent behaviour in these time and temperature histories?
4. Are all devices similar in the same frequency range?
5. Do correlations exist between sensitivity and conductance?

In addition to investigating these questions using the parallel coordinate technique, a set of statistical analyses have been undertaken in order to give further credence to any possible relationships that the visualisations suggest. There are a number of other interesting features of the data that can be illustrated clearly using parallel coordinates, and these features are also described in this case study.

Questions 1 to 3: the effects of temperature and time

In order to investigate the relationship between the measurements and the water temperature, it was necessary to look at a single device at a single frequency. If more than one device or more than one frequency was examined, the plot became too confused to interpret. Figure 24 shows an example of the plots used.

The data shown in figure 24 was taken from the measurements made over the second frequency range, 15 to 85 kHz. The axes represent device number, date, month, temperature, and conductance and sensitivity at 15, 50, and 85 kHz. The axes have been reordered to put the temperature axis next to the sensitivity at 15 kHz axis, and the sensitivity at 15 kHz axis has been reversed. The data set has been focussed to show only selected data points, namely all measurements made using device 4. The line segments between the 5th and 6th axes, labelled Temperature and Sens_15, do not cross one another very often, and when they do cross the angles are comparatively flat (compare for example the angles between the Date and Month axes). This lack of intersection means that there is a near-monotonic inverse relationship between temperature and sensitivity at 15 kHz for device number 4.

A number of other similar relationships were discovered for other transducers at other frequencies. Some relationships were inverse and some were direct. The relationships identified are listed in table 1. All of these discoveries agreed with the statistical analyses.

Dev. no.	Freq. range (kHz)	Freq. (kHz)	Measurement	Direct or inverse
1	15–85	15	Sensitivity	Inverse
3	15–85	15	Sensitivity	Inverse
4	15–85	15	Sensitivity	Inverse
1	15–85	50	Sensitivity	Direct
3	15–85	50	Sensitivity	Direct
4	15–85	50	Sensitivity	Direct
4	70–220	70	Sensitivity	Inverse
2	200–350	200	Sensitivity	Inverse
2	200–350	200	Conductance	Inverse
1	200–350	350	Sensitivity	Direct

Table 1: Relationships identified between temperature and measurement data.

Similar relationships have been identified between the measurement values and time. An example plot is shown in figure 25. The data have been taken from all measurements made on device 4 in the 1 – 20 kHz range. There are few line segment intersections between the Date and Cond_1 axes, indicating that at 1 kHz there is a monotonic relationship between date and conductance measurement. The same is true of the lines between the Date and Cond_8 axes. The monotonic nature of these relationships means that they may well be caused by drift: any seasonal affects such as temperature dependence would be more likely to have maxima and minima.

Other similar relationships have been identified, and are listed in full in table 2. These relationships agreed with the statistical analyses of the data. All were direct monotonic relationships, making drift a possibility. Drift is particularly likely for devices 1 – 4 in the 1 – 20 kHz range and device 3 in the 15 – 85 kHz range, since these transducers exhibited monotonic relationships at all three frequencies examined.

Further investigation is needed to answer the question of how the frequency affects the dependencies of conductance and sensitivity on temperature and time, although the tables above give some indication. More intermediate frequencies would have to be examined, and as yet this has not been done.

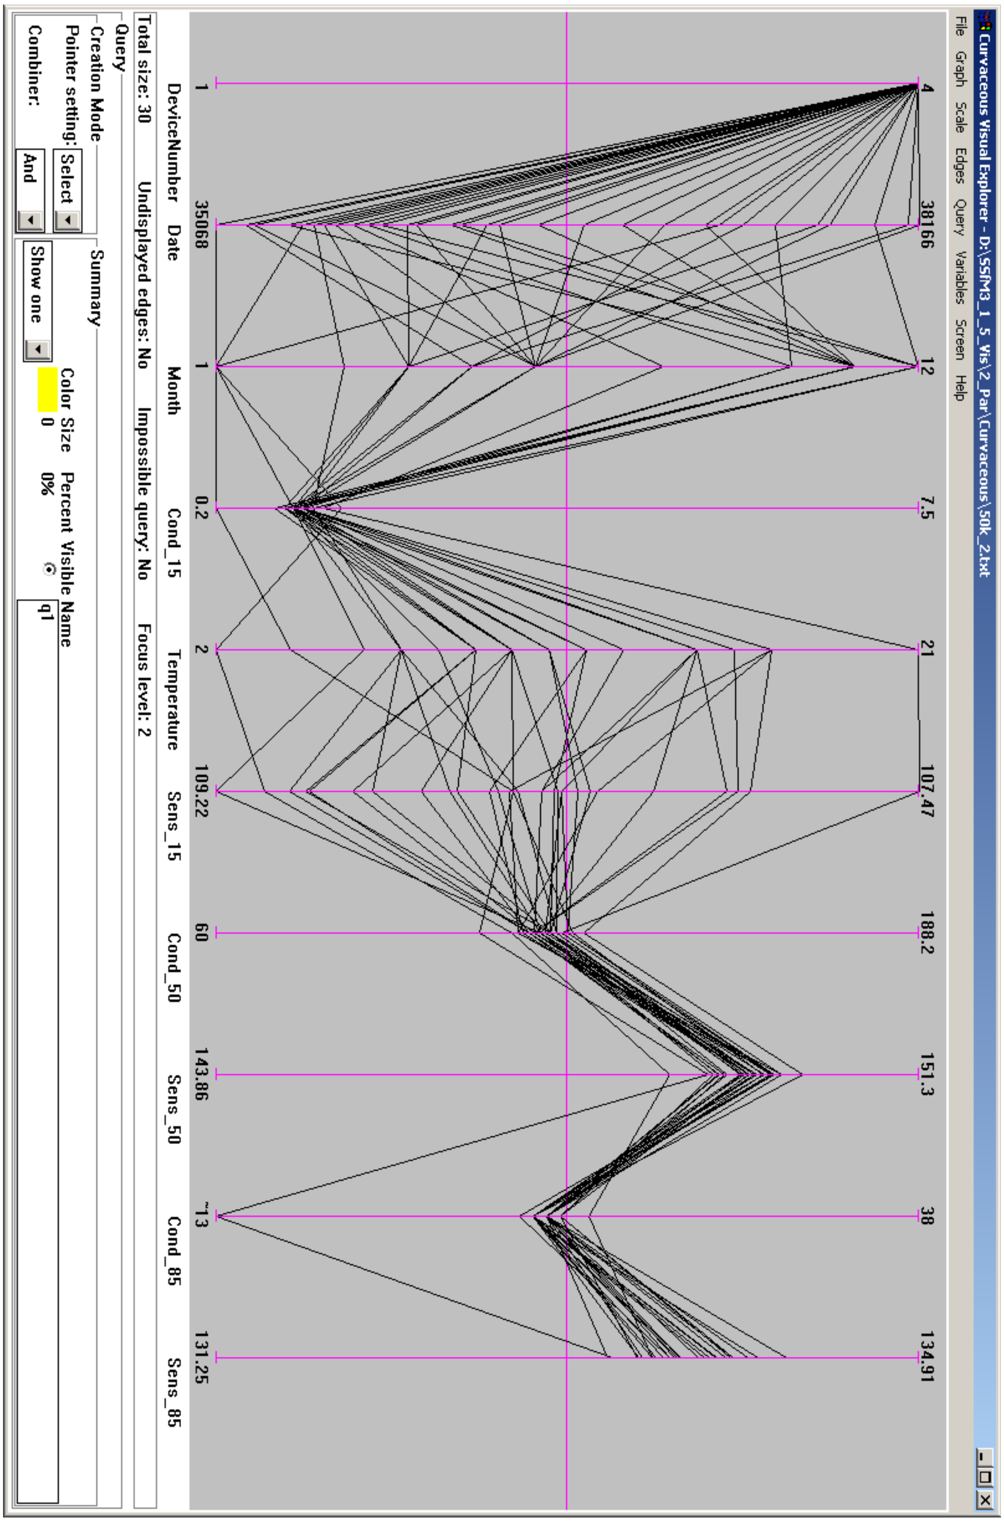

Figure 24: Measured data from device number 4 in the 15–85 kHz range. Displayed data indicates an inverse near-monotonic relationship between temperature and sensitivity at 15 kHz.

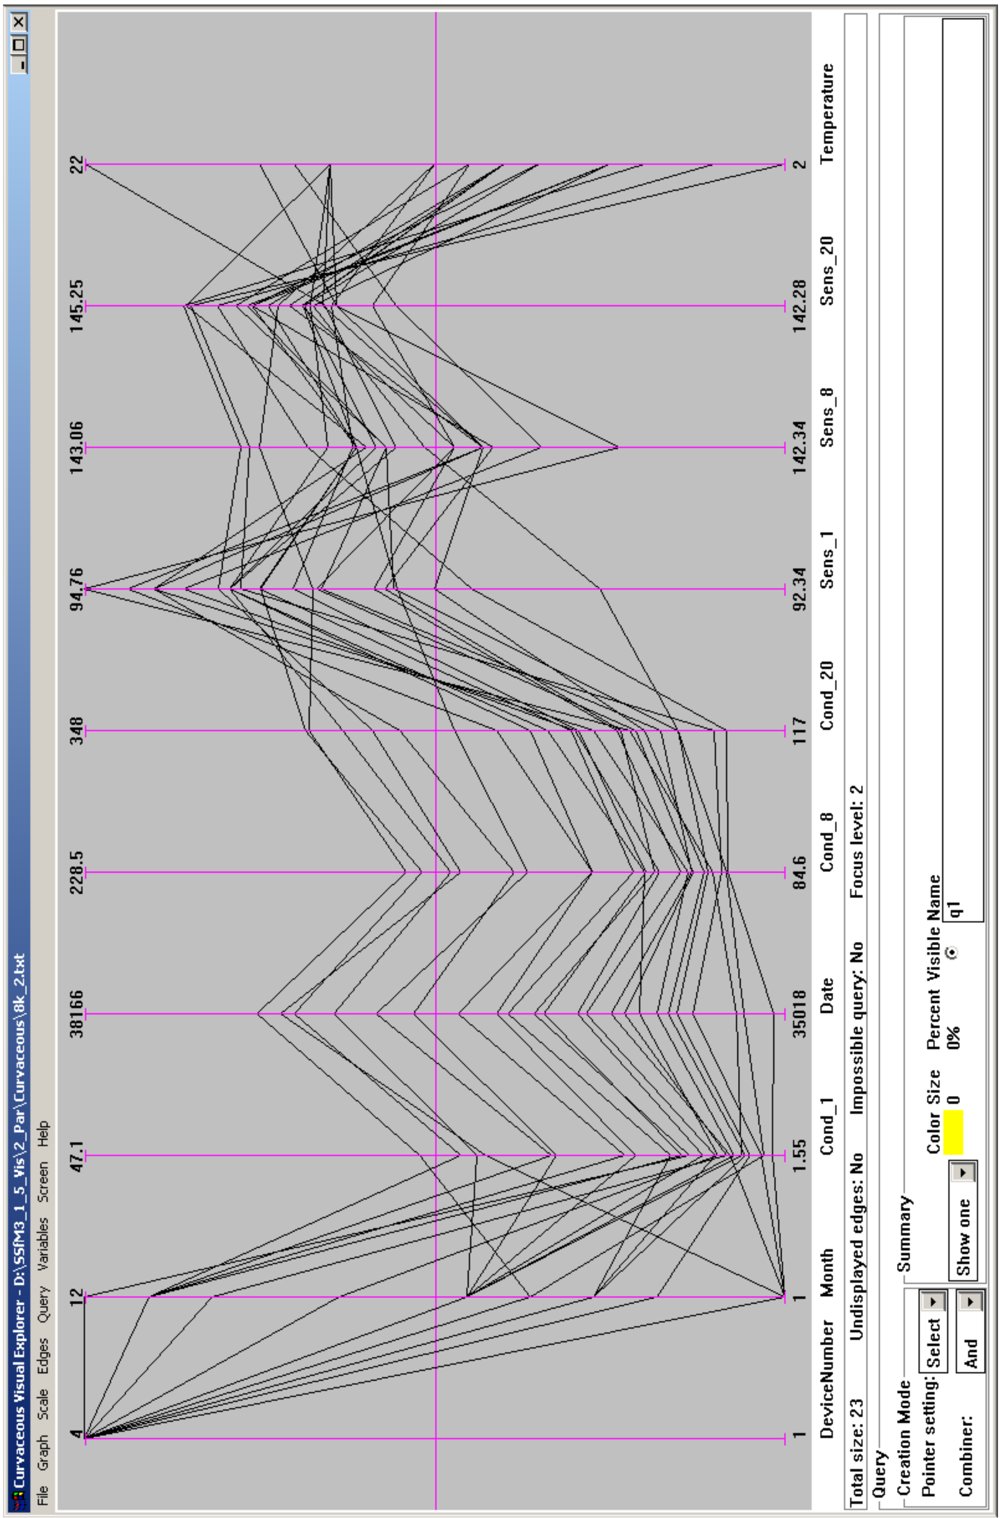

Figure 25: Measured data from device number 4 in the 1–20 kHz range. Displayed data indicates a direct near-monotonic relationship between date and conductance at 1 and 8 kHz.

Dev. no.	Freq. range (kHz)	Freq. (kHz)	Measurement	Direct or inverse
1–4	1–20	1	Conductance	Direct
1–4	1–20	8	Conductance	Direct
1–4	1–20	20	Conductance	Direct
3	15–85	15	Conductance	Direct
4	15–85	15	Conductance	Direct
3	15–85	50	Conductance	Direct
3	15–85	85	Conductance	Direct
1	70–220	150	Conductance	Direct

Table 2: Relationships identified between time and measurement data.

Question 4: Similarity of devices

Different devices were compared by plotting the measurement values for each device on a separate axis, then highlighting a date of interest to look for any relationships. An example is shown in figure 26. The data shown are all measured values from the transducers with a frequency range of 1 – 20 kHz. The highlighted data are the results measured on 4th December 1999, a date when all four transducers were measured. The axes have been ordered so that the four conductivity axes and the four sensitivity axes are consecutive so that similar performance can be identified.

It is clear from figure 26 that the line segments connecting the sensitivity axes are parallel and nearly horizontal, and so the sensitivities of the different devices are linearly related. Examination of the limits on the axes shows that the maxima and minima are almost the same, so it is likely that the sensitivities of the various devices are almost the same.

The conductance data has a more complicated structure. At the lower end of the conductance scale, there is a near-monotonic relationship between the conductances of the different devices, but near the higher end of the conductance axis, which corresponds to higher frequencies, the line segments between axes cross over one another, meaning that the devices behave differently. Even if the axes are reordered there is no clear relationship between any of the pairs of devices.

The other frequency ranges do not exhibit such clear-cut behaviour. Figure 27 shows data from the transducers with a frequency range of 15 – 85 kHz. There was no date on which all of the transducers were measured, so two separate days with similar temperatures have been chosen, and the other data have been removed to improve ease of interpretation. This choice means that some axes have missing data points (if no measurement was made for a transducer on a date), which are represented as measurement values at the bottom of the relevant axes. Hence all sets of line segments starting at the bottom of an axis and fanning out along an adjacent axis should be ignored.

If the misleading line segments are ignored, it is clear from the non-crossing line segments that the sensitivities of different devices have a monotonic and near-linear relationship for

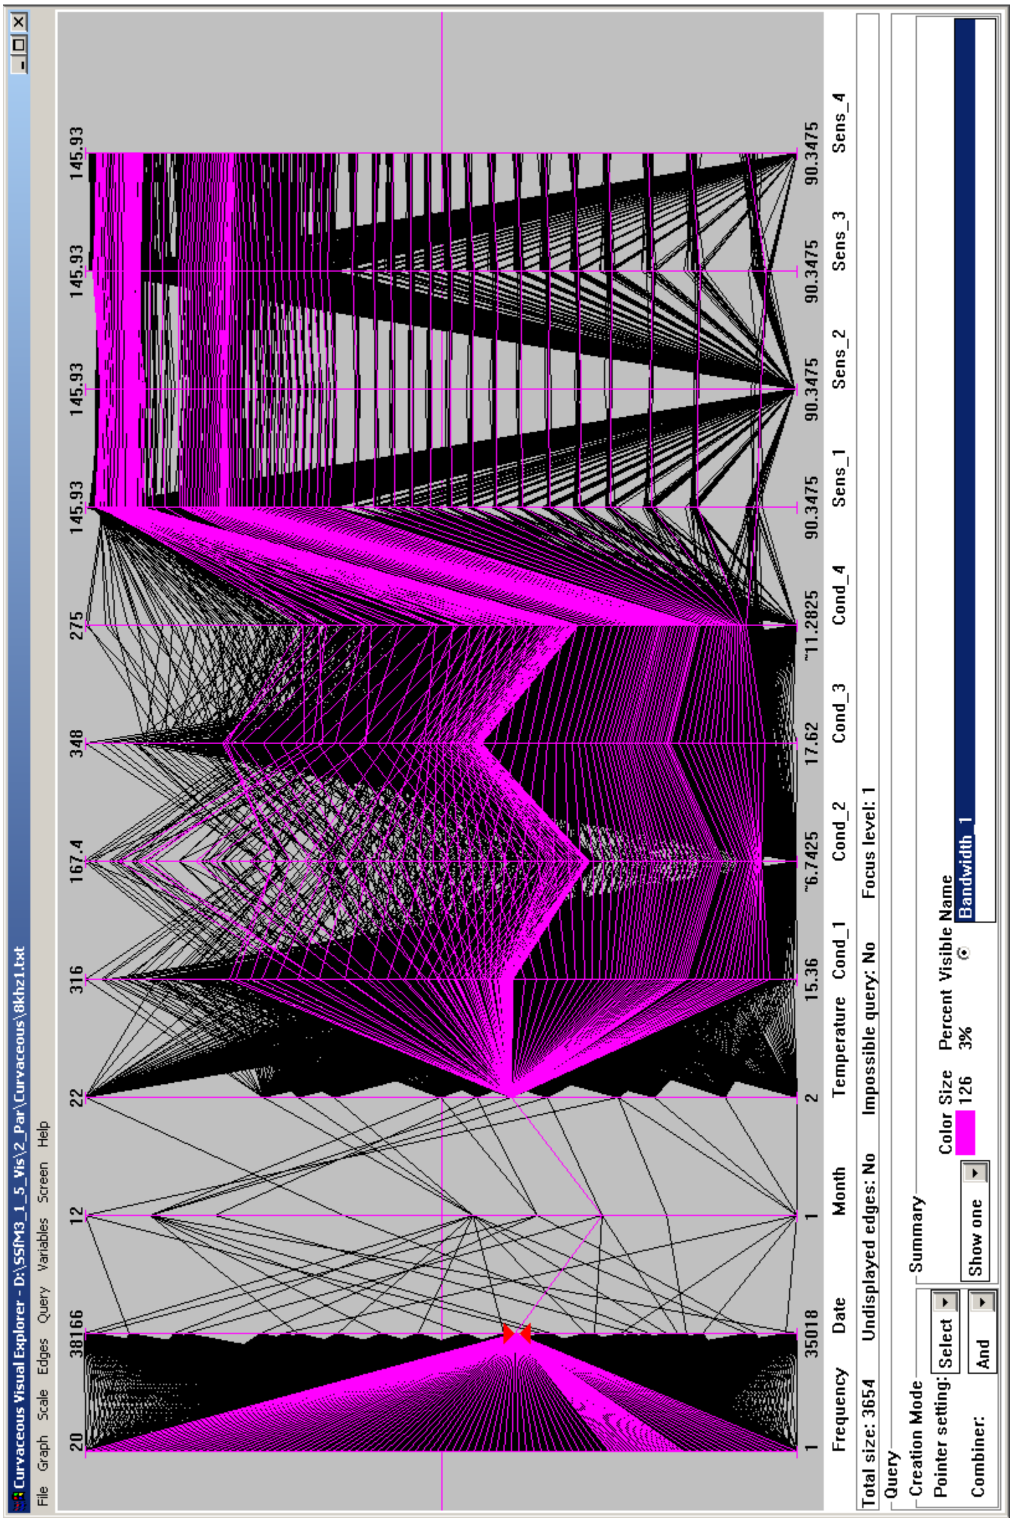

Figure 26: Data from a single date at 1 – 20 kHz from all devices, showing similar sensitivity behaviour for all devices.

low values of sensitivity, but that the relationship becomes more complicated at higher values of sensitivity. The line segments between the Cond_1 and Cond_4 axes are nearly parallel and rarely cross, meaning that the conductance behaviours of transducers 1 and 4 are close to linearly related.

The 70 – 220 kHz and 200 – 350 kHz ranges look similar to this plot: sensitivities are nearly linearly related for low values of sensitivity, but the relationship becomes more complex at higher values, and the conductance behaviour varies from device to device.

Question 5: Correlations between sensitivity and conductance

Figure 28 shows data collected on a single date from all transducers in the 1 – 20 kHz range. The line segments between the conductance and sensitivity axes corresponding to the same device do not cross very much at low values of conductance and sensitivity (corresponding to low frequencies), indicating a monotonic relationship between the two. At higher values of conductance and sensitivity the line segments cross over and the relationship between conductance and sensitivity becomes more complicated.

Plots generated using other frequency ranges and other dates have similar properties: for low values of sensitivity and conductance, there is a monotonic relationship between the two measurands, but for higher values the relationship becomes more complicated.

Other points of interest

Several other points of interest were noticed during the data investigation. Figure 29 shows data from all the devices with a 1 – 20 kHz range. Measurements made at 1, 8, and 20 kHz are shown, with the three conductance axes adjacent. The small number of crossing line segments between the conductance axes, and the small angles at which the line segments cross, indicate that for these devices the conductance measurements at 1, 8, and 20 kHz are strongly correlated. No other similar behaviour was noted in any other frequency ranges or for any sensitivity measurements. No other frequencies within the 1 – 20 kHz range have been examined.

Some evidence of unusual transducer behaviour has been found. Figure 30 shows one such event. The plot shows measurements made at 200, 300, and 350 kHz. The highlighted point is a single set of measurements made using device 1. All three conductance measurements are abnormally high, suggesting that there may have been a problem with the conductance measurements of that device on that day. No such events were found in the data from the other devices on that date.

Figure 31 was generated by plotting the full set of data in the 200 – 350 kHz range, and highlighting line segments that pass through the top 16% of the conductance axis. Tracking the line segments back to the first axis indicates that all of these points were generated using

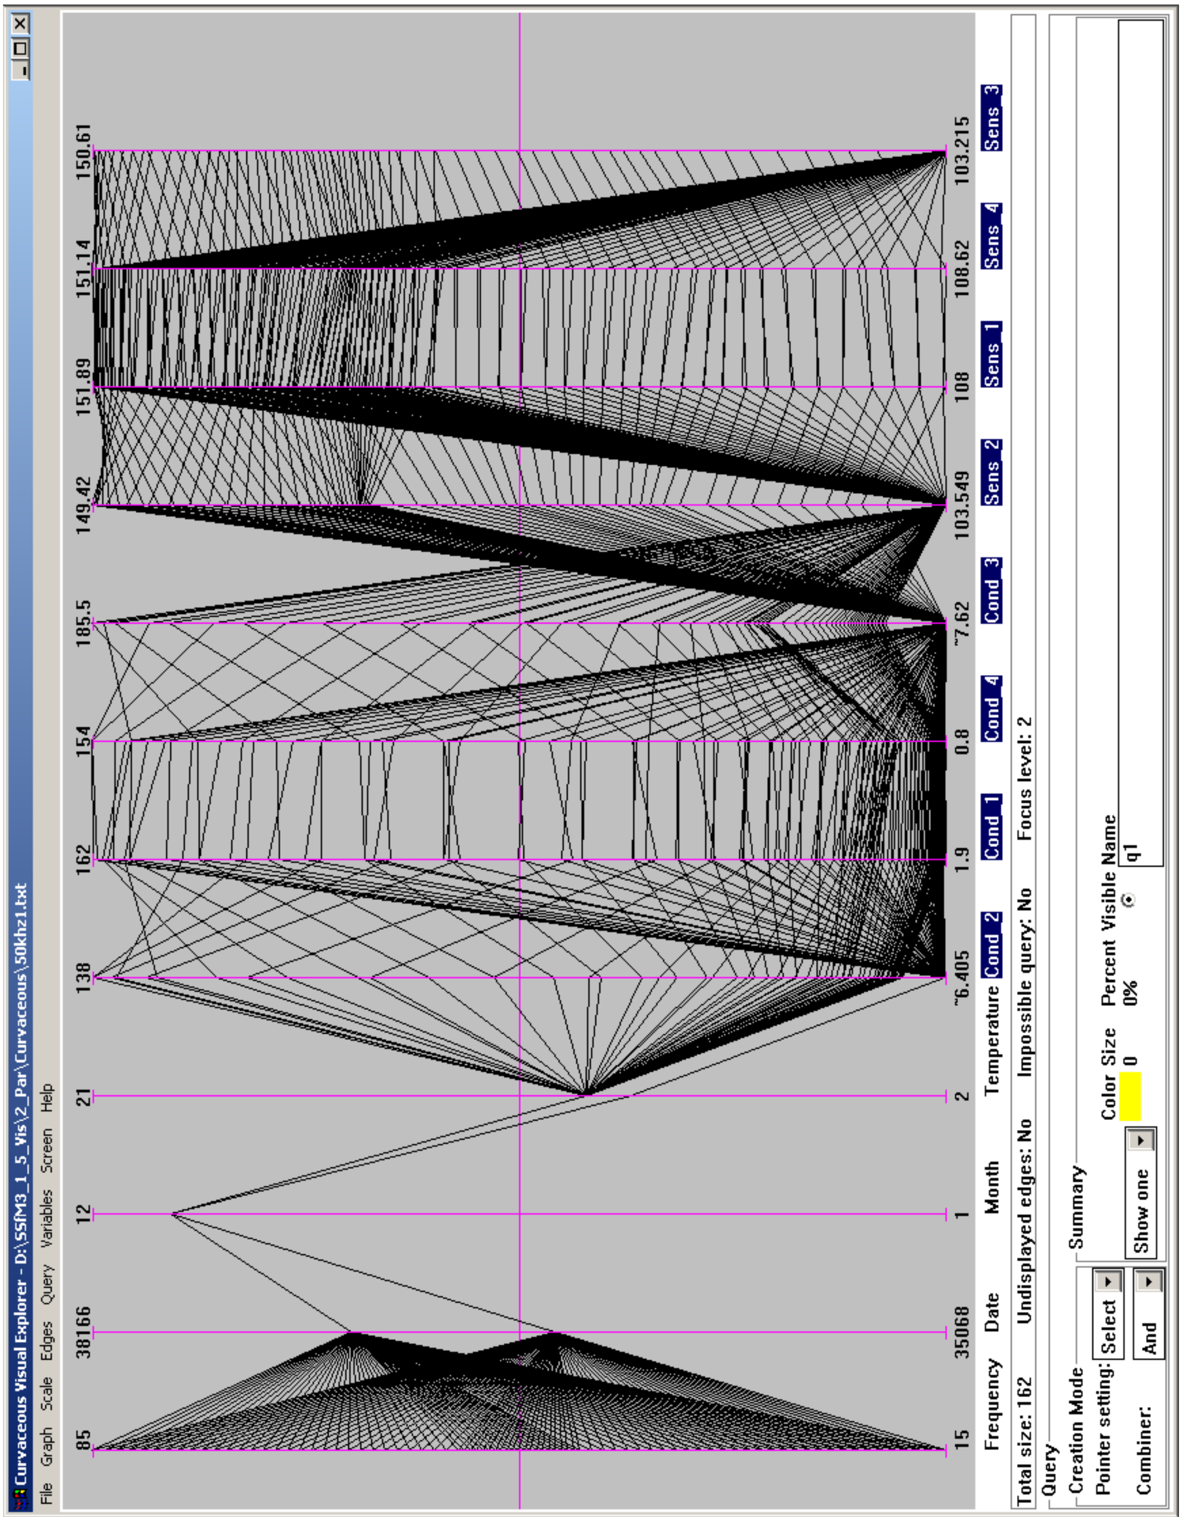

Figure 27: Data from two dates with similar temperatures at 15 – 85 kHz from all devices.

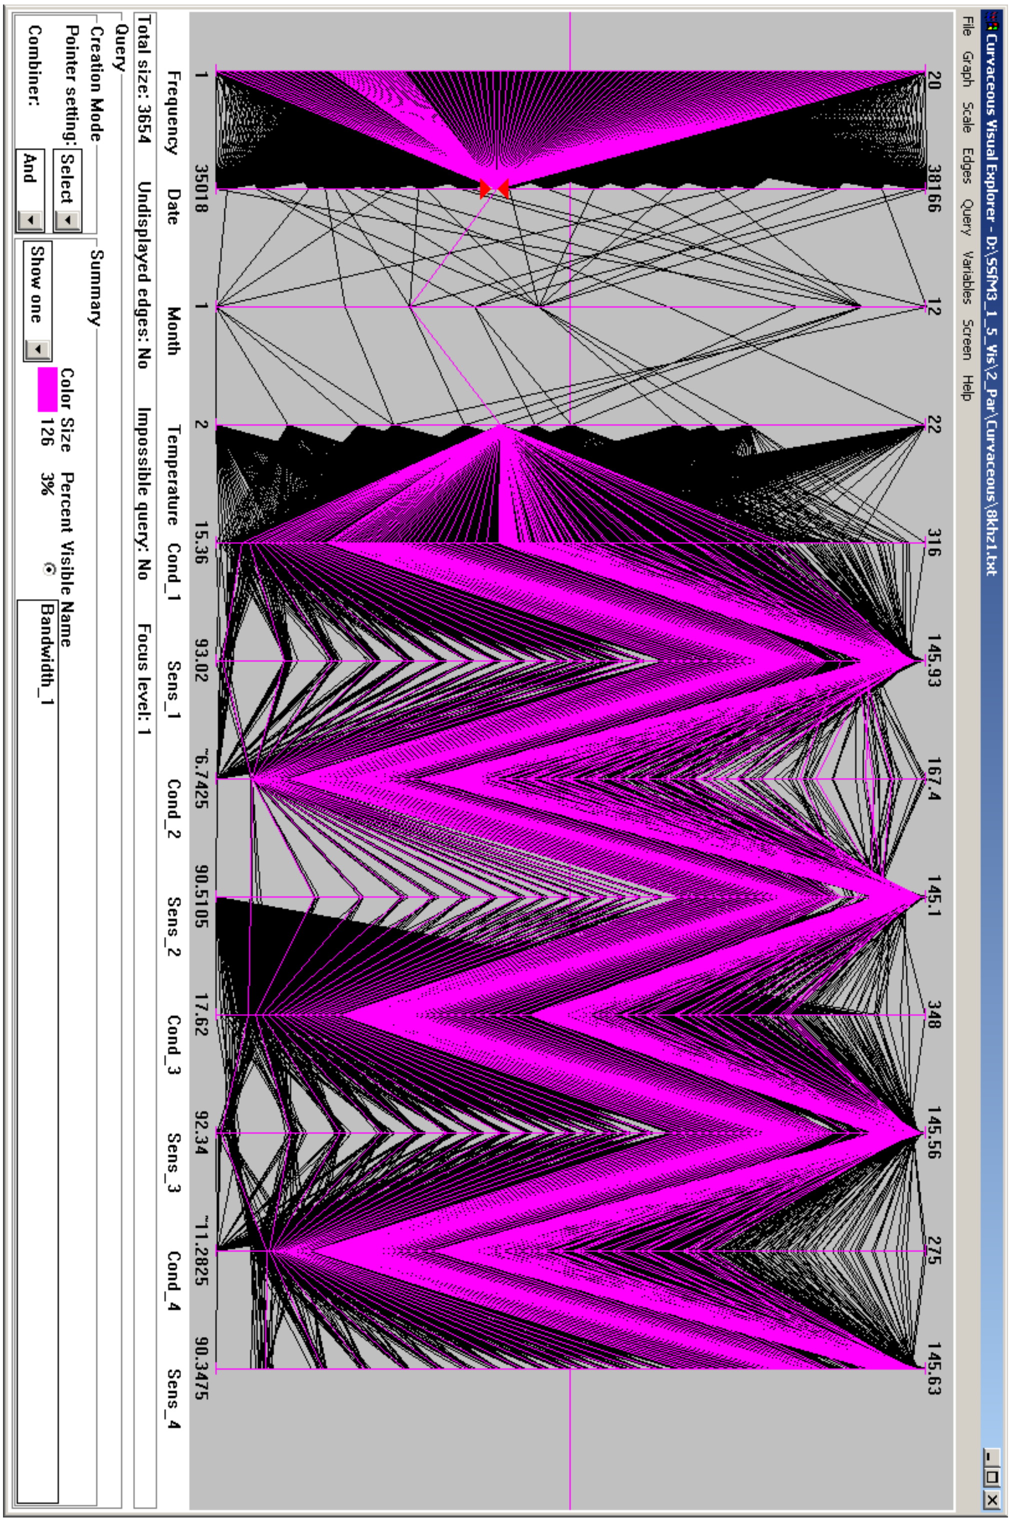

Figure 28: Data from a single date from all transducers in the 1 – 20 kHz range. Highlighted lines show a possible monotonic relationship between conductance and sensitivity at low frequencies.

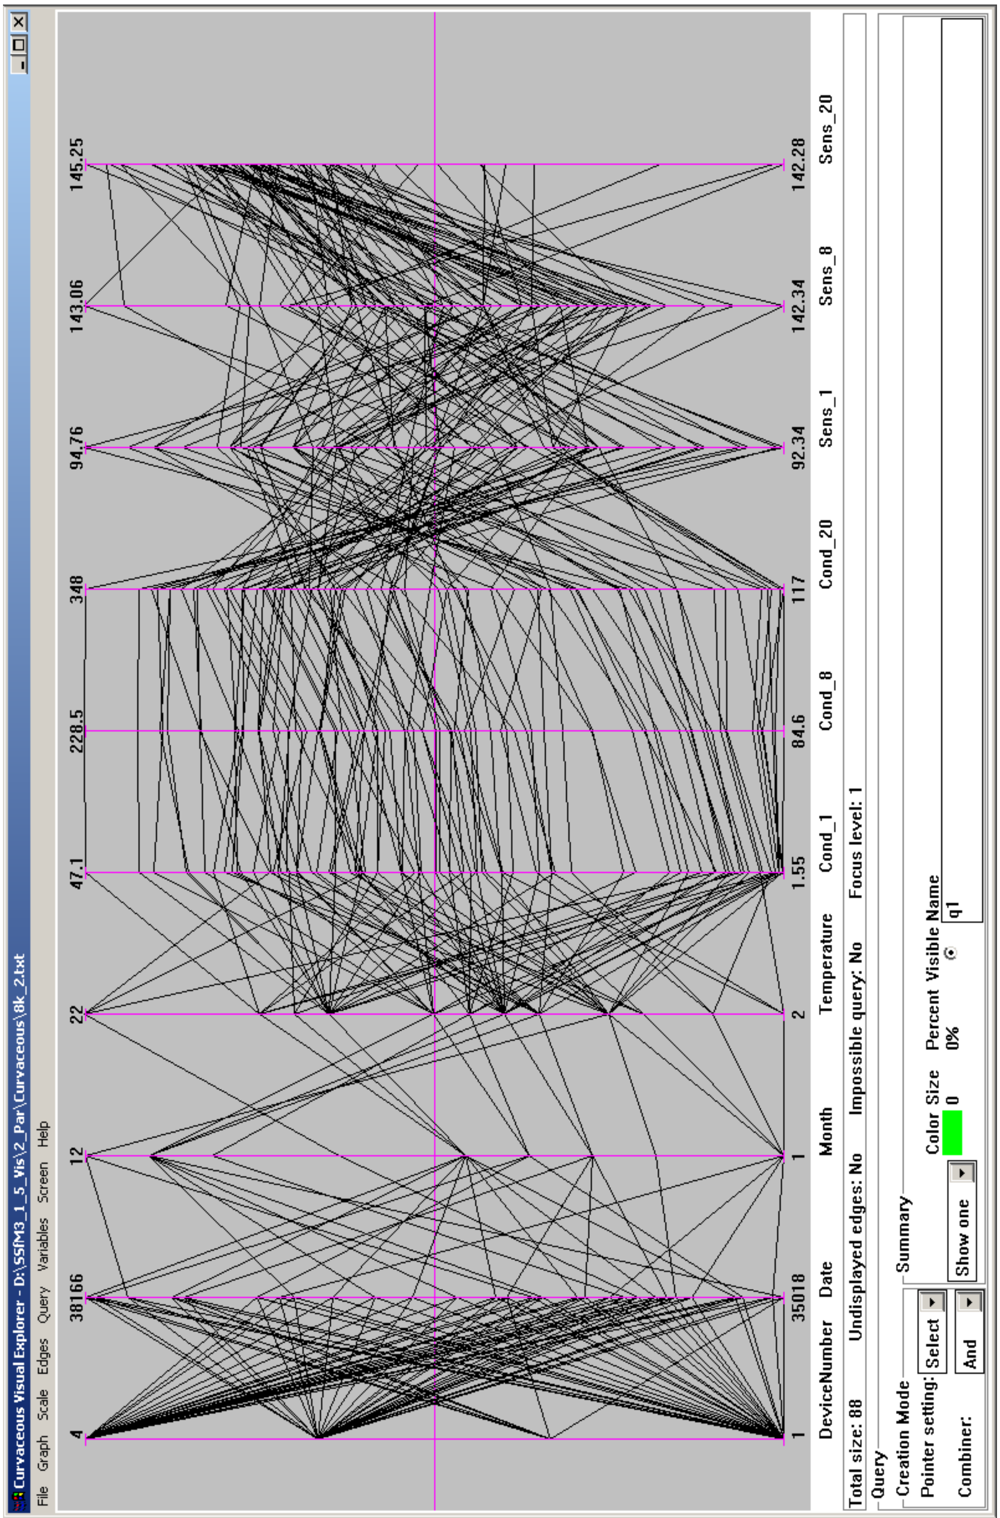

Figure 29: Measurements made at 1, 8, and 20 kHz for all devices. The conductance axes show evidence of correlation between the measurements from each device for different frequencies.

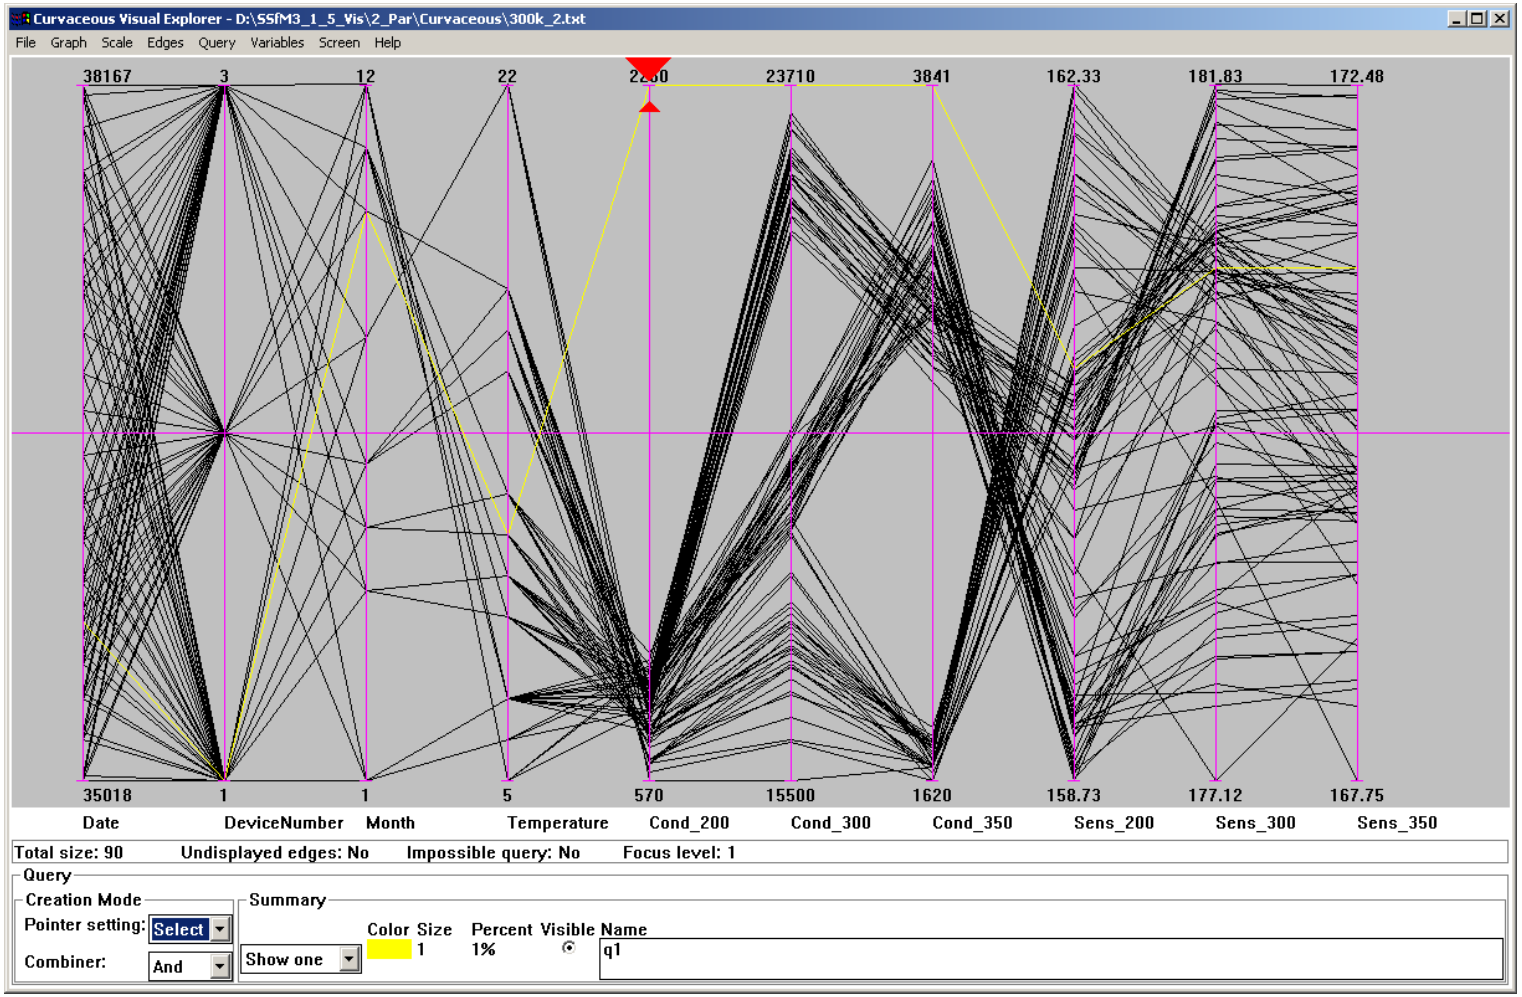

Figure 30: Data in the 200 – 350 kHz band. The highlighted point shows unusually high capacitance measurements in a single device.

device 1, meaning that device 1 has the highest peak conductance in this range.

Figures 32 and 33 compare the behaviours of devices 1 and 2 in the 1 – 20 kHz range. The interesting point is that the range of conductance values measured from device 1 is much larger than the range of conductance values measured from device 2. It is clear from the date axis that device 1 has been measured on more occasions than device 2, but the ranges of sensitivity values from the two devices are about the same, meaning that the lack of data is unlikely to be what has caused the limited conductance range. Devices 3 and 4 of the same frequency range are similar to device 1, indicating that device 2 is exhibiting unusual behaviour.

All of the conductance measurement values should be positive. Figure 34 highlights some measurements that produced a negative conductance value. The measurements were all taken from a single device on a single occasion, and so are likely to indicate a one-off problem. Any other plots showing negative conductance values have been generated in examples containing missing data. If a data value is missing, CVE automatically represents it as the lower limit of the axis minus 5% of the range of the axis, which can lead to negative values for positive quantitites.

5.1.4 Conclusions

The use of the parallel coordinate technique made it possible to visualise a large multivariate data set quickly and easily. Interpretation of the plots has led to some useful information about device behaviour, and some unusual events and behaviours have been identified by using the technique.

5.2 Case Study 2: An industrial drying problem

5.2.1 The data set

The data used in this case study comes from an industrial application in tile drying. Modifications were made to a clay tile dryer, and the data sets were gathered in order to investigate whether the modifications had made any difference to the tile-drying process. The aim of the drying process is to remove most of the water from the tiles before they are put into the kiln for firing. The tiles need to be dried to the same degree as quickly as possible, but if they are dried too fast they can crack or warp.

Tiles are stacked on large cars to be dried. Each car consists of about 10 trays stacked above one another with a small amount of air space between them, each tray typically holding about 10 – 20 tiles. The cars are placed in a large room called a chamber dryer. The chamber dryer is gradually heated, and hot air is blown across the tiles by fans in order to make the conditions within the dryer as uniform as possible so that all the tiles dry at the

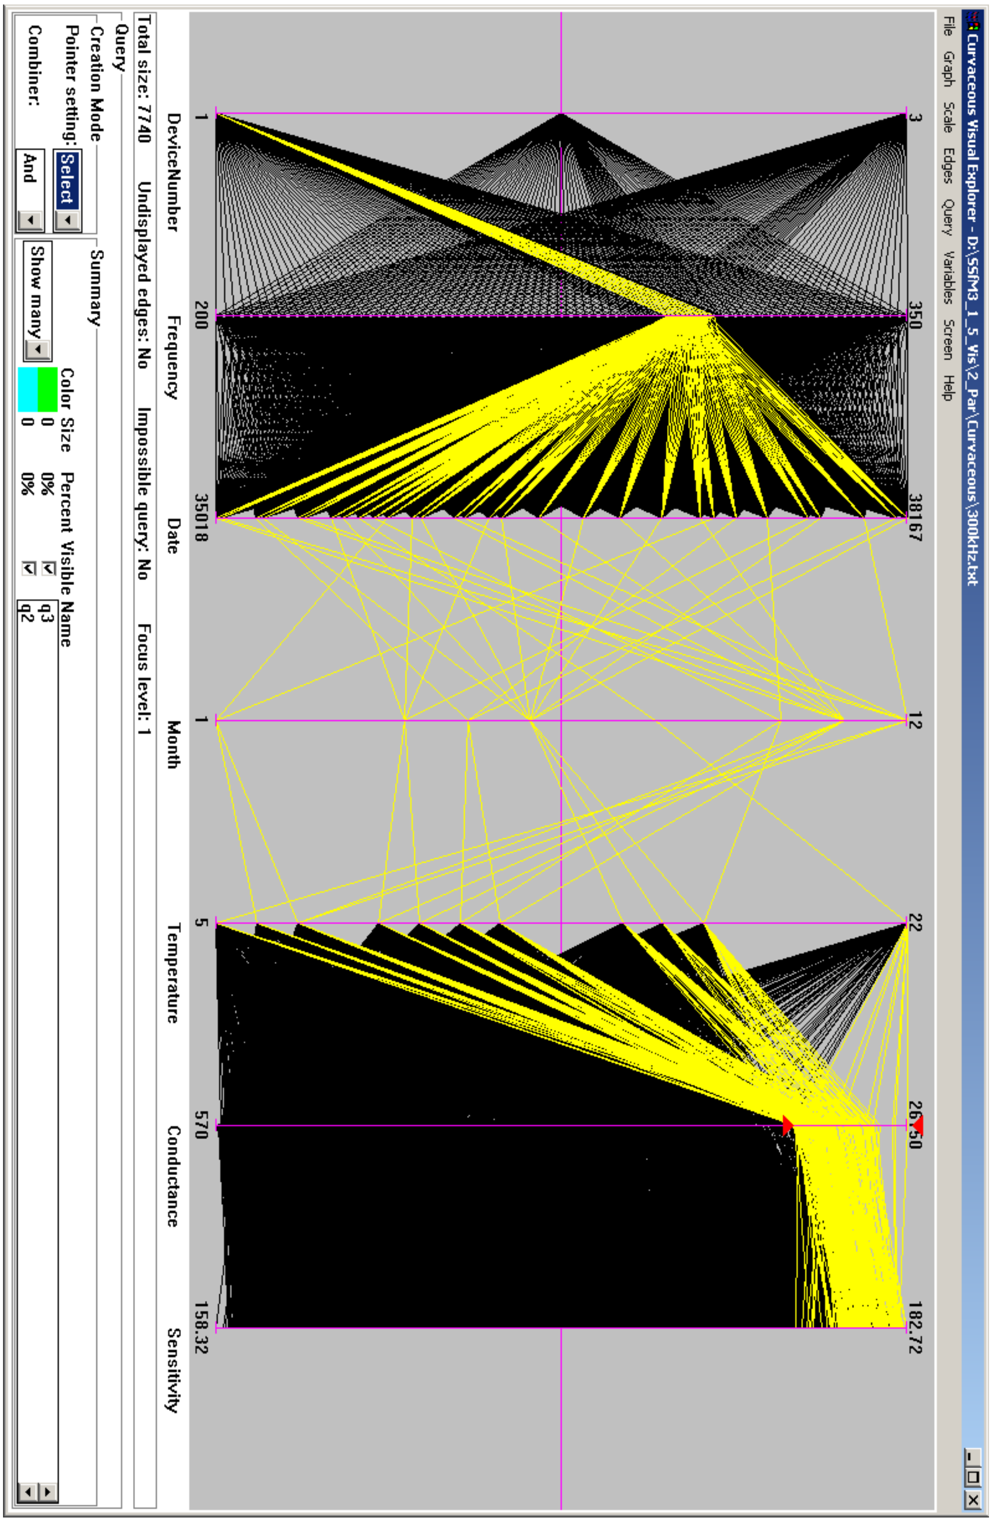

Figure 31: Data in the 200 – 350 kHz band. The highlighted points lie on the top 16% of the conductance axis.

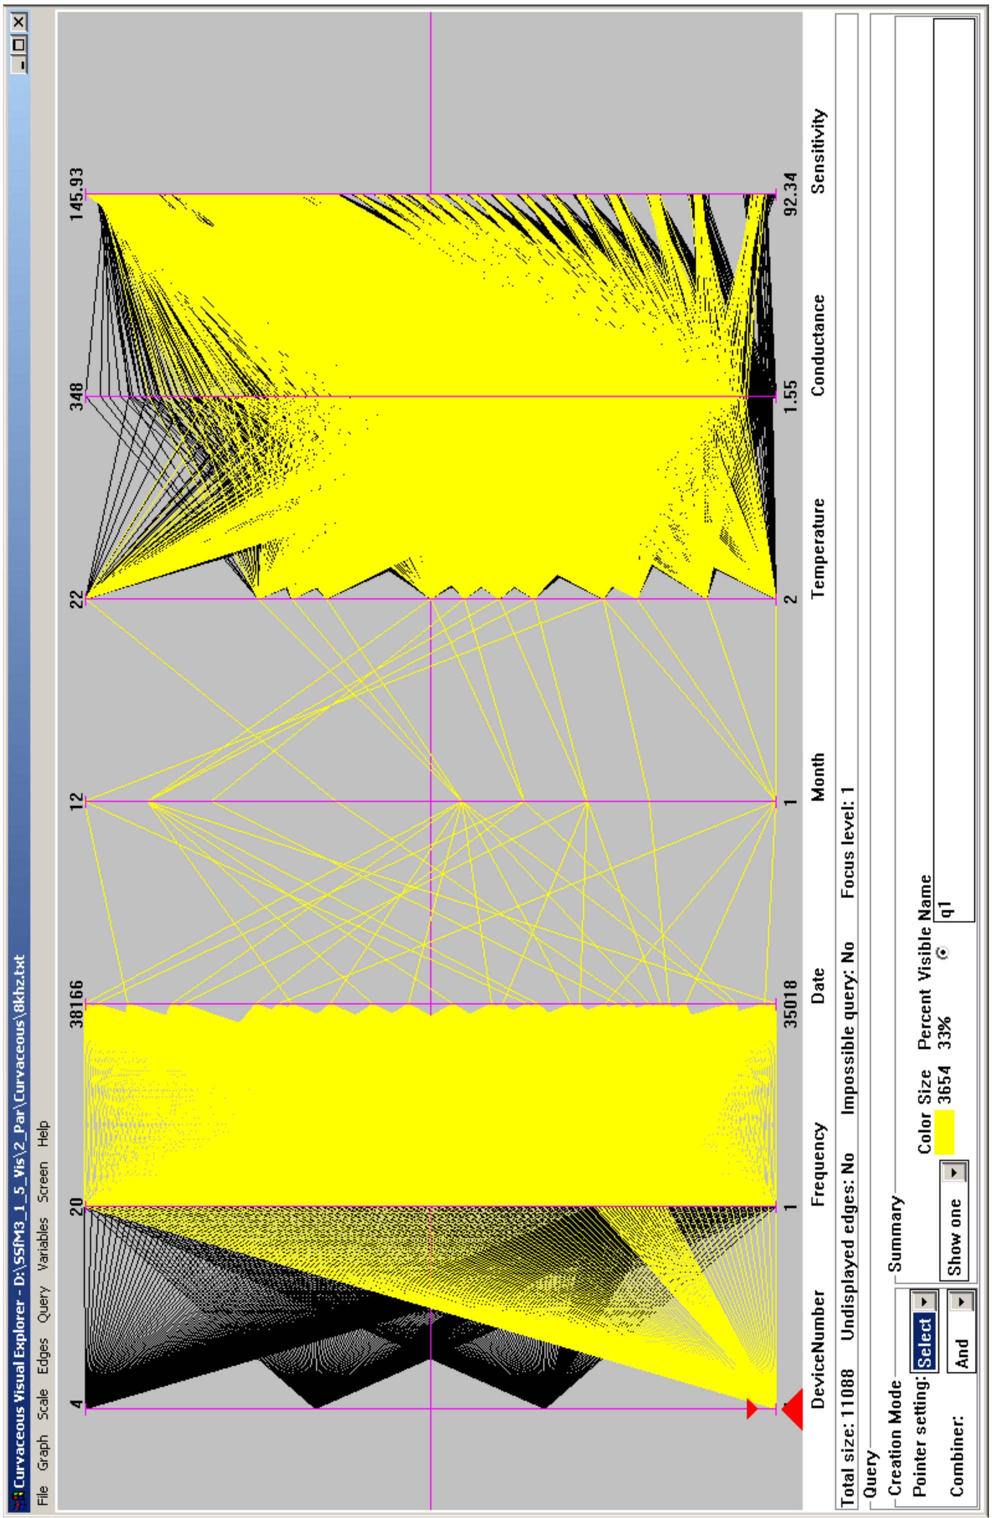

Figure 32: Data in the 1 – 20 kHz band. The highlighted points were measured using device 1.

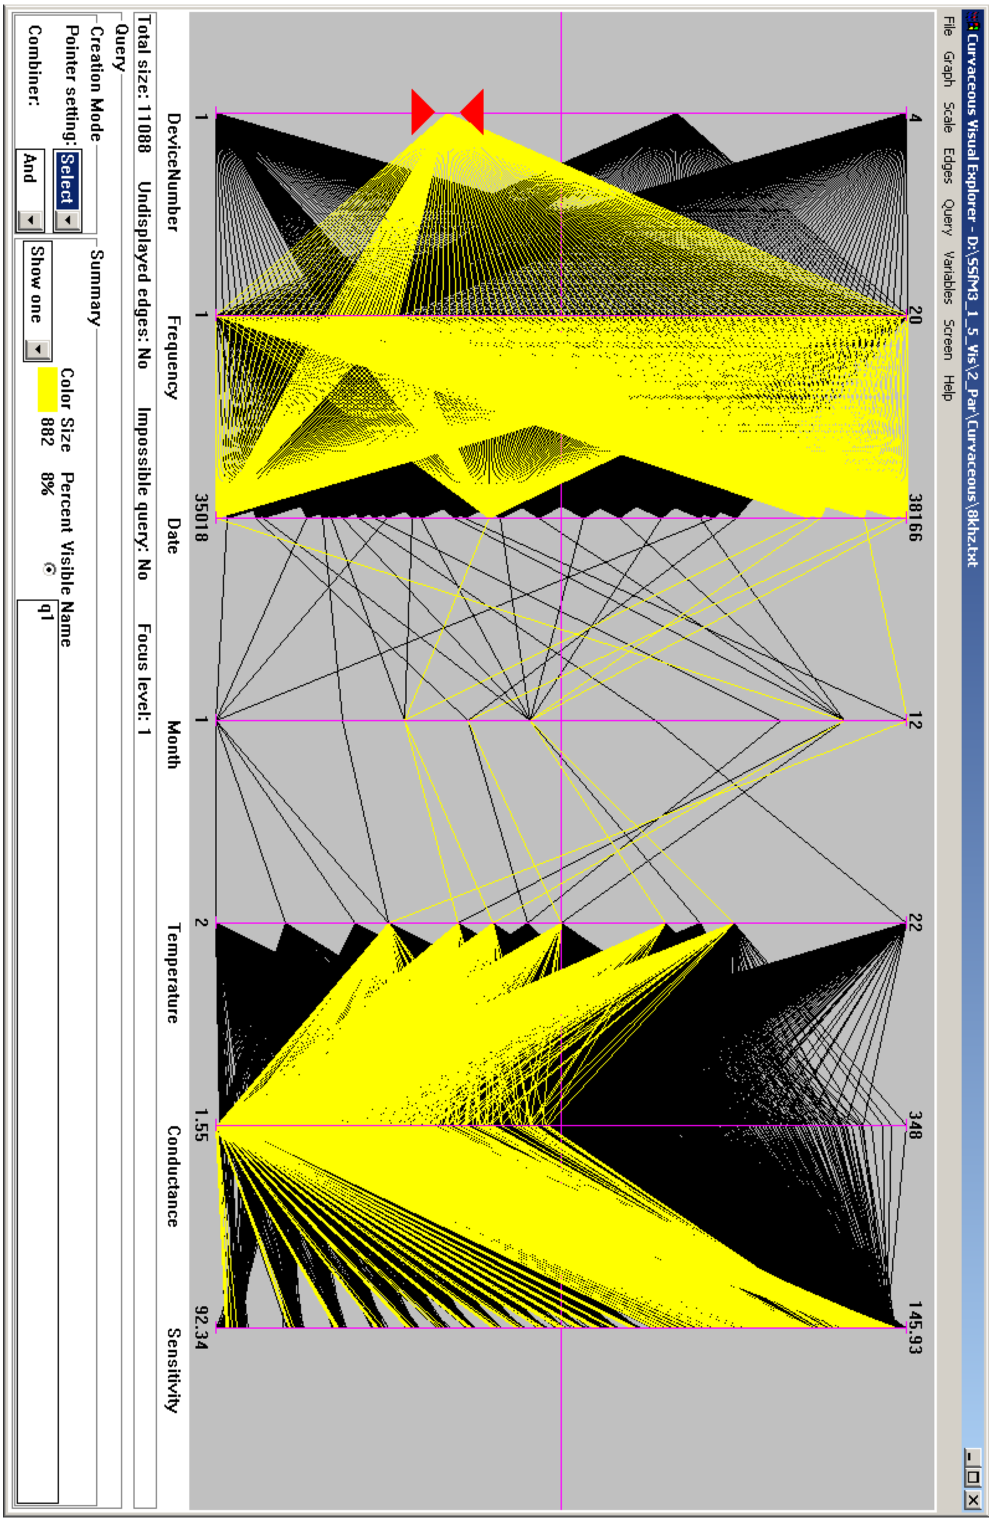

Figure 33: Data in the 1 – 20 kHz band. The highlighted points were measured using device 2.

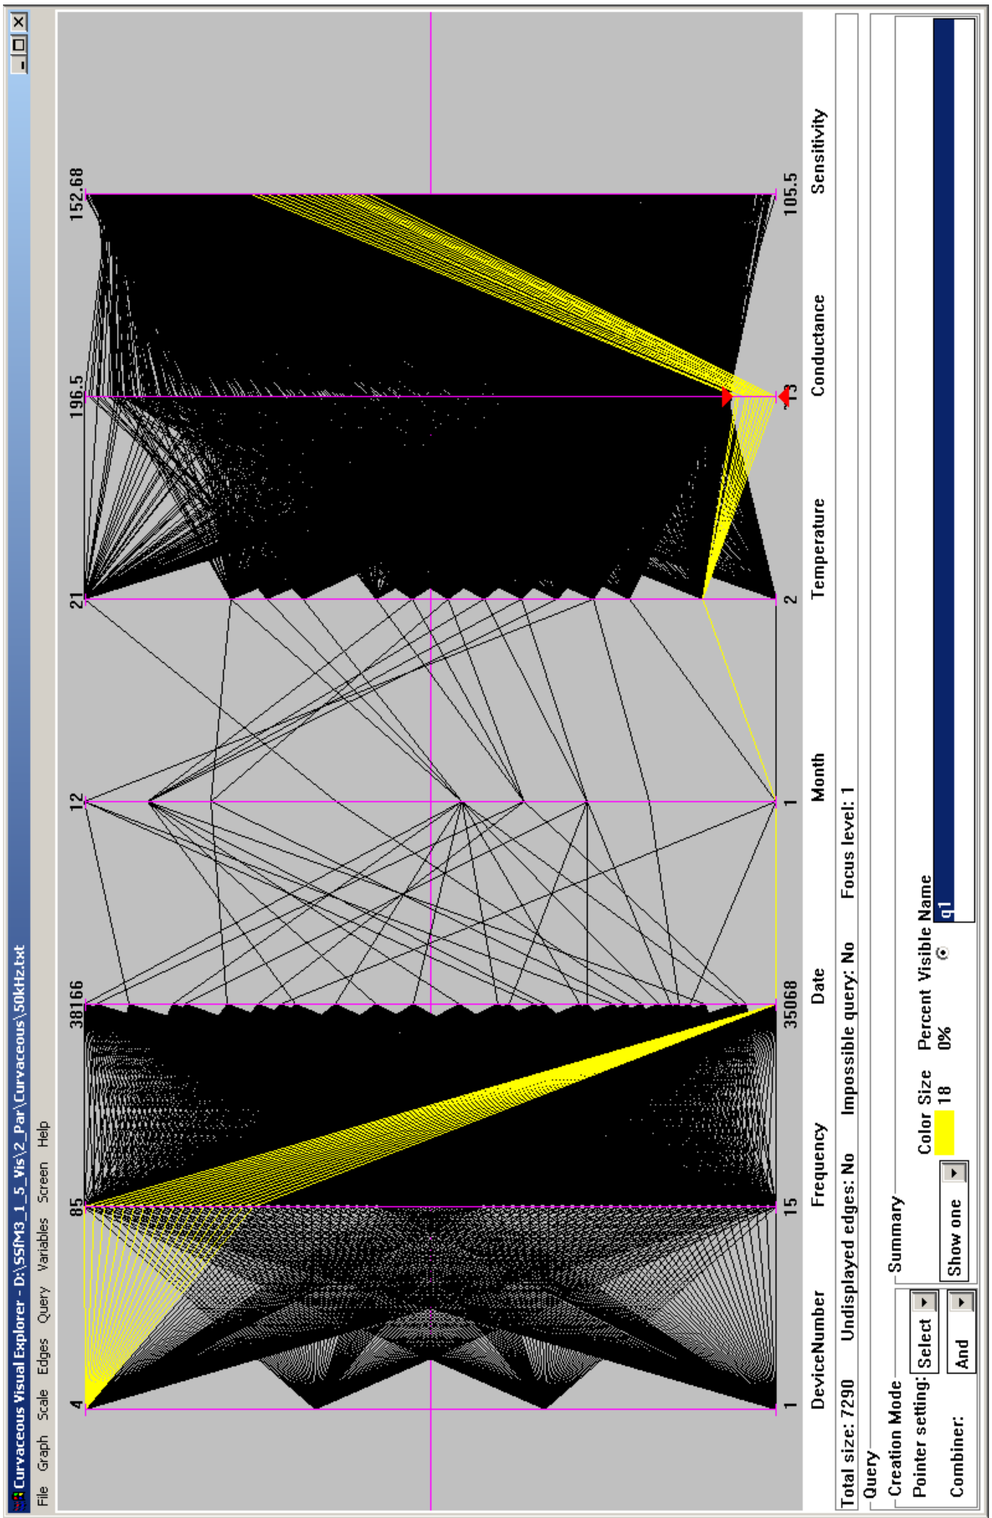

Figure 34: Data in the 15 – 85 kHz band. The highlighted points show negative conductance values.

same speed. The drying process generally takes tens of hours, and in this case the tiles were dried for fifty hours.

Drying is driven by two processes: the evaporation of water from the surface of the tile, and the diffusion of water within the tile to feed the evaporation. The evaporation depends on the velocity of the air across the tile, and the temperature and humidity of the air. The diffusion within the tile depends on the clay structure and the tile temperature. Tile temperature is also increased by heat transfer from other sources such as thermal radiation and conduction from the metal supports, which are at the bulk air temperature rather than the evaporation-cooled tile surface temperature.

As the tile dries, it shrinks because water leaves the structure. Shrinkage is easier to measure than weight loss without disrupting the drying process, and so shrinkage is used as a substitute for dryness for data logging purposes. It is assumed that the two are directly (linearly) equivalent. The equivalence is reasonable since thermal expansion is unlikely to have any effect at the temperatures used for drying, and the linearity assumptions is likely to be reasonable over the range of tile moisture contents seen in drying.

The data sets that were taken measured the shrinkage of nine different tiles during the drying process, and measured the air temperature, the wet bulb temperature, and the tile temperature at two different locations within a drying car. The wet bulb temperature is the temperature of a water-saturated object in air and is dependent on the humidity of the air and the air temperature. It is measured in order to calculate humidity. A set of values was measured every five minutes for the whole of the fifty hour drying cycle. When the measurements were repeated after the modifications had been made, the same sensors were used to take measurements in each location. Figure 35 shows the location of the different sensors: an S indicates the location of a shrinkage sensor and TRH indicates the location of a set of temperature sensors.

5.2.2 The nature of the data

Temperatures were measured using 8-bit loggers with a resolution of 0.1 °C, and shrinkage sensors with a resolution of approximately 0.01%. The low resolution of the data makes calculation of shrinkage rates very difficult as wide moving average ranges are needed to smooth the data and a shrinkage rate calculation may only give four discrete rate values.

Additionally, under some circumstances the shrinkage data can behave oddly in the initial stages of drying: the sensor may not move at all, or it may move quite far quite quickly, both of which are physically unlikely. The first is likely to be caused by plastic deformation around the sensor, and the second may be caused by settling problems. In the first case the drying rate will be underestimated, leading to loss of data, and in the second case it will be over-estimated, but usually this can be spotted. It is generally expected that the drying rate will be slow initially because the dryer is still heating up, slow finally because there is little water left in the tile, and faster in the time between.

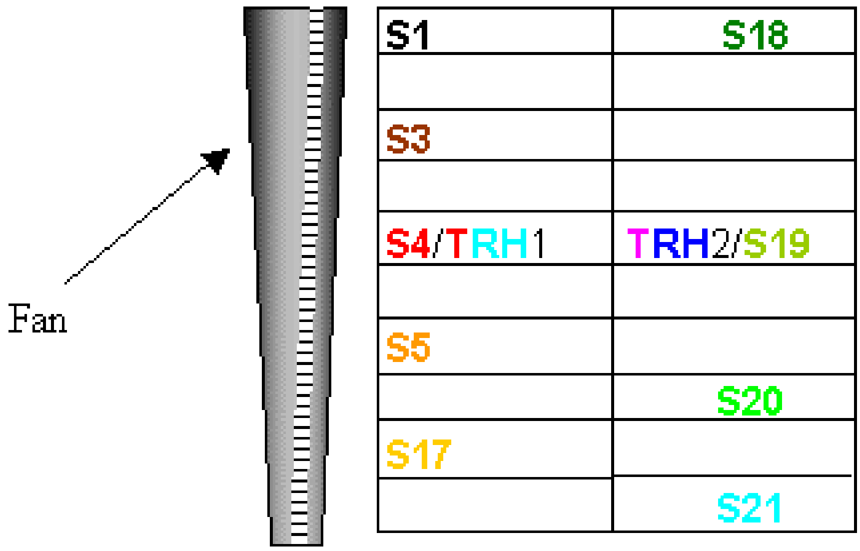

Figure 35: Measurement positions within a dryer car relative to a fan.

5.2.3 Questions to be answered by the visualisations

The visualisation process aimed to answer three questions:

1. Is the drying rate dependent on position of the tile within the car, and if so what correlations can be seen between tiles in different places?
2. Did the modifications to the dryer affect the temperature conditions and the drying rates?
3. Is there any clear relationship between the various temperature measurements and the corresponding tile shrinkages and drying rates?

The visualisations shown in this section were generated using Curvaceous Software's C:Visual Explorer (see section 4 for details). Points of interest are mostly indicated by coloured line segments. Red triangles are used to manipulate the interactive highlighting tools, and are often indicating the limits of a selected range of values. The significance of the coloured line segments is generally explained in the accompanying text.

Question 1: Effects of position

The effects of position on the temperature conditions and shrinkage rates can be determined by comparing sets of data that were gathered at the same horizontal level (i.e. sets 1 and 18 and sets 4 and 19), and sets that were gathered at the same vertical level (i.e. sets 1, 3, 4, 5, and 17, and sets 18, 19, 20, and 21). These four sets of data have been plotted before and after the modifications to the dryer using the parallel coordinate technique.

Figure 36 plots all of the shrinkage and shrinkage rate data gathered before modifications were made on a set of parallel axes. Each axis (other than the initial time axis) is labelled with the sensor position number and Rate or Shrink depending on whether it is shrinkage data or shrinkage rate data. The highlighted points indicate the highest values of the shrinkage rate for tiles 4 (yellow lines, fastest 11% highlighted) and 19 (light blue lines, fastest 16% highlighted). The two sets of lines are separate and mutually exclusive. Looking at the time axis shows that the fastest shrinkage of tile 19 occurs before the fastest shrinkage of tile 4, implying that the ambient conditions (temperature, humidity, and air velocity) around the two tiles are unlikely to be the same.

Figure 37 shows the same data gathered after the modifications to the dryer. The high shrinkage rates of tile 4 are highlighted in green (fastest 15% highlighted), and those of tile 19 in pink (fastest 11% highlighted). The two sets of points are no longer mutually exclusive (this is shown by the presence of pink lines between the red triangles on the axis labelled Rate 4), and the fastest shrinkage rates of the two tiles occur at roughly the same time. This similarity indicates that the modifications have made the ambient conditions around the two tiles more similar, possibly by improving the air flow so that better air mixing occurred.

Figures 38 and 39 show the same information for tiles 1 (green line segments in figure 38, 11% highlighted, yellow line segments in figure 39, 10% highlighted) and 18 (pink line segments in figure 38, 13% highlighted, light blue line segments in figure 39, 15% highlighted). The regions of fastest drying overlap before and after the dryer modifications. This overlap suggests that the ambient conditions around these tiles were similar before the modifications. This is probably because the tiles were on the top layer of the dryer car, which is open to the atmosphere and so has a free air flow above it. It is also interesting to note that the fast shrinkage of tiles 1 and 18 occurs before the fast shrinkage of tiles 4 and 19, suggesting that the free air flow enhances the drying rate. This would be expected since good air mixing and higher air velocity result in faster drying.

Figure 40 shows the data collected before modifications to the dryer, and the top 17% shrinkage rates for tile 4 are highlighted in light blue. The axes have been ordered so as to make comparison between the shrinkage rates of tiles 1, 3, 4, 5, and 17 more straightforward. The highlighted points show that in general if tile 4 is shrinking quickly then so are tiles 3, 5, and 17, but tile 1 is not. Looking at the time axis and the axis labelled Shrink 1, and considering the information gained from figure 38, it can be seen that by the time tile 4 is drying rapidly, tile 1 has already gone through a period of rapid shrinkage and is quite dry. These features of the figures indicate that the top-to-bottom uniformity of the side of the dryer car nearest to the fan is quite good, but that tile 1 dries more quickly than the other tiles (probably because it is open to the air).

Figure 41 shows the same data collected after modifications to the dryer, with the top 10% shrinkage rates for tile 4 are highlighted in green and the top 10% shrinkage rates for tile 1

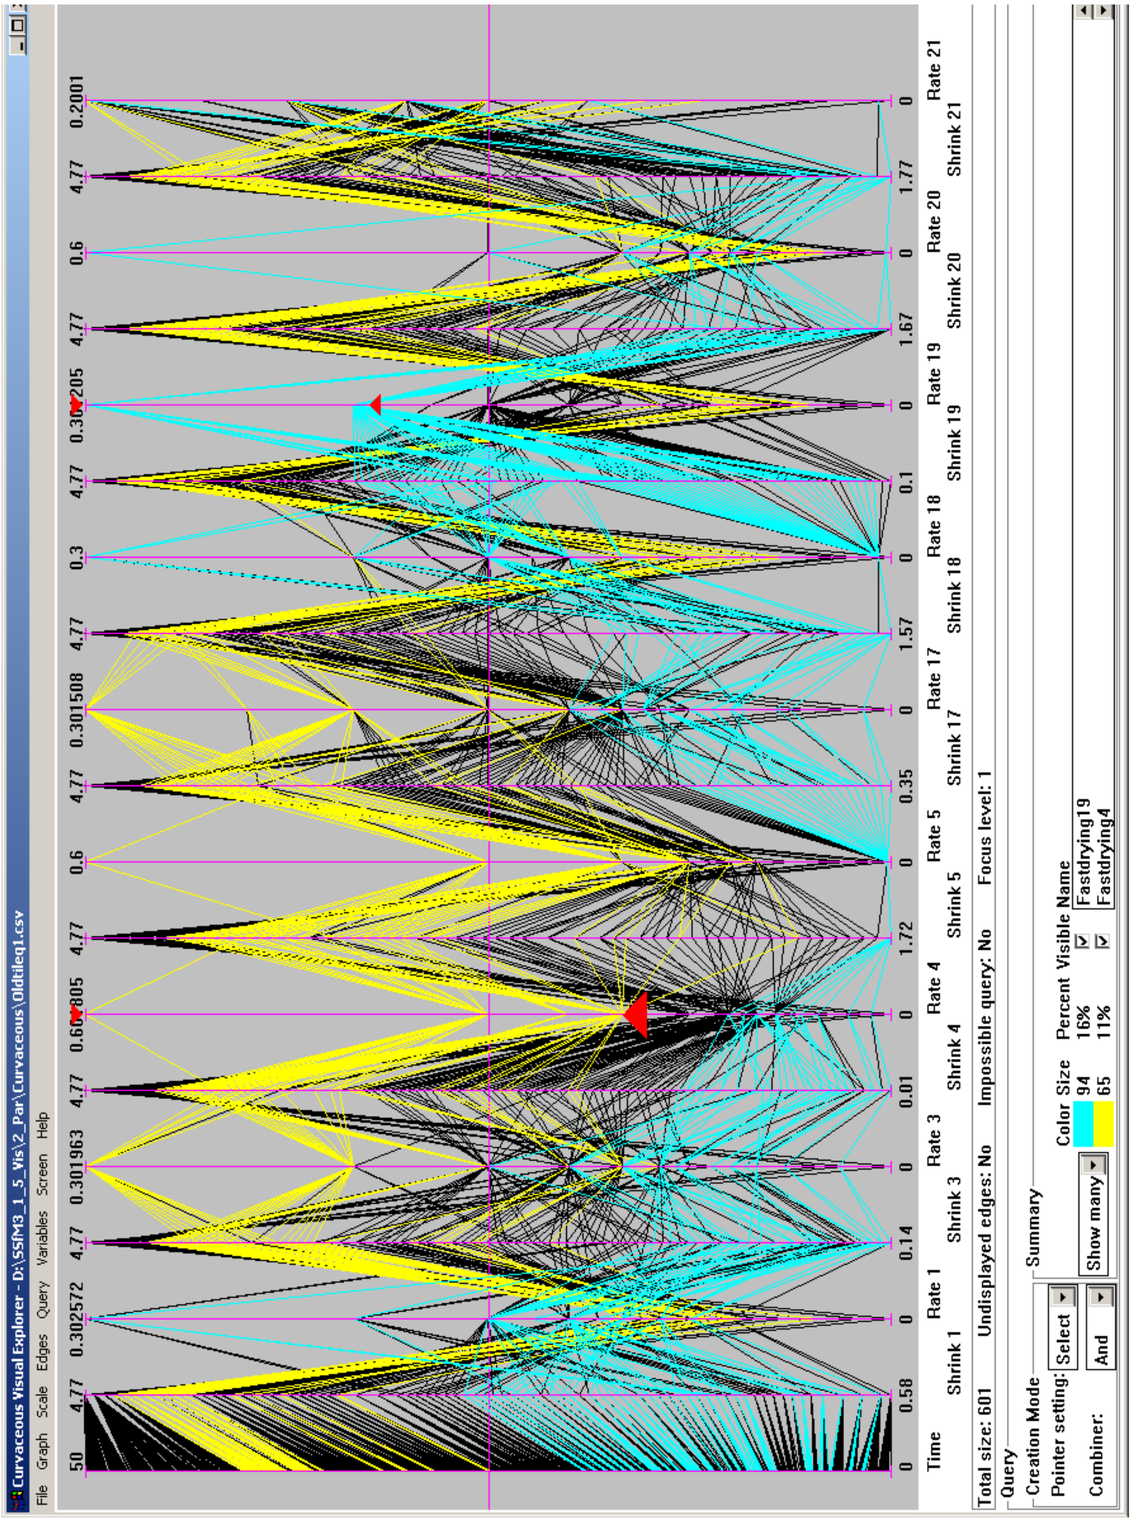

Figure 36: Data gathered before modifications to the dryer. Highlighted points are the fastest shrinkage rates for tiles 4 (yellow) and 19 (light blue). The fastest shrinkages occur at different times.

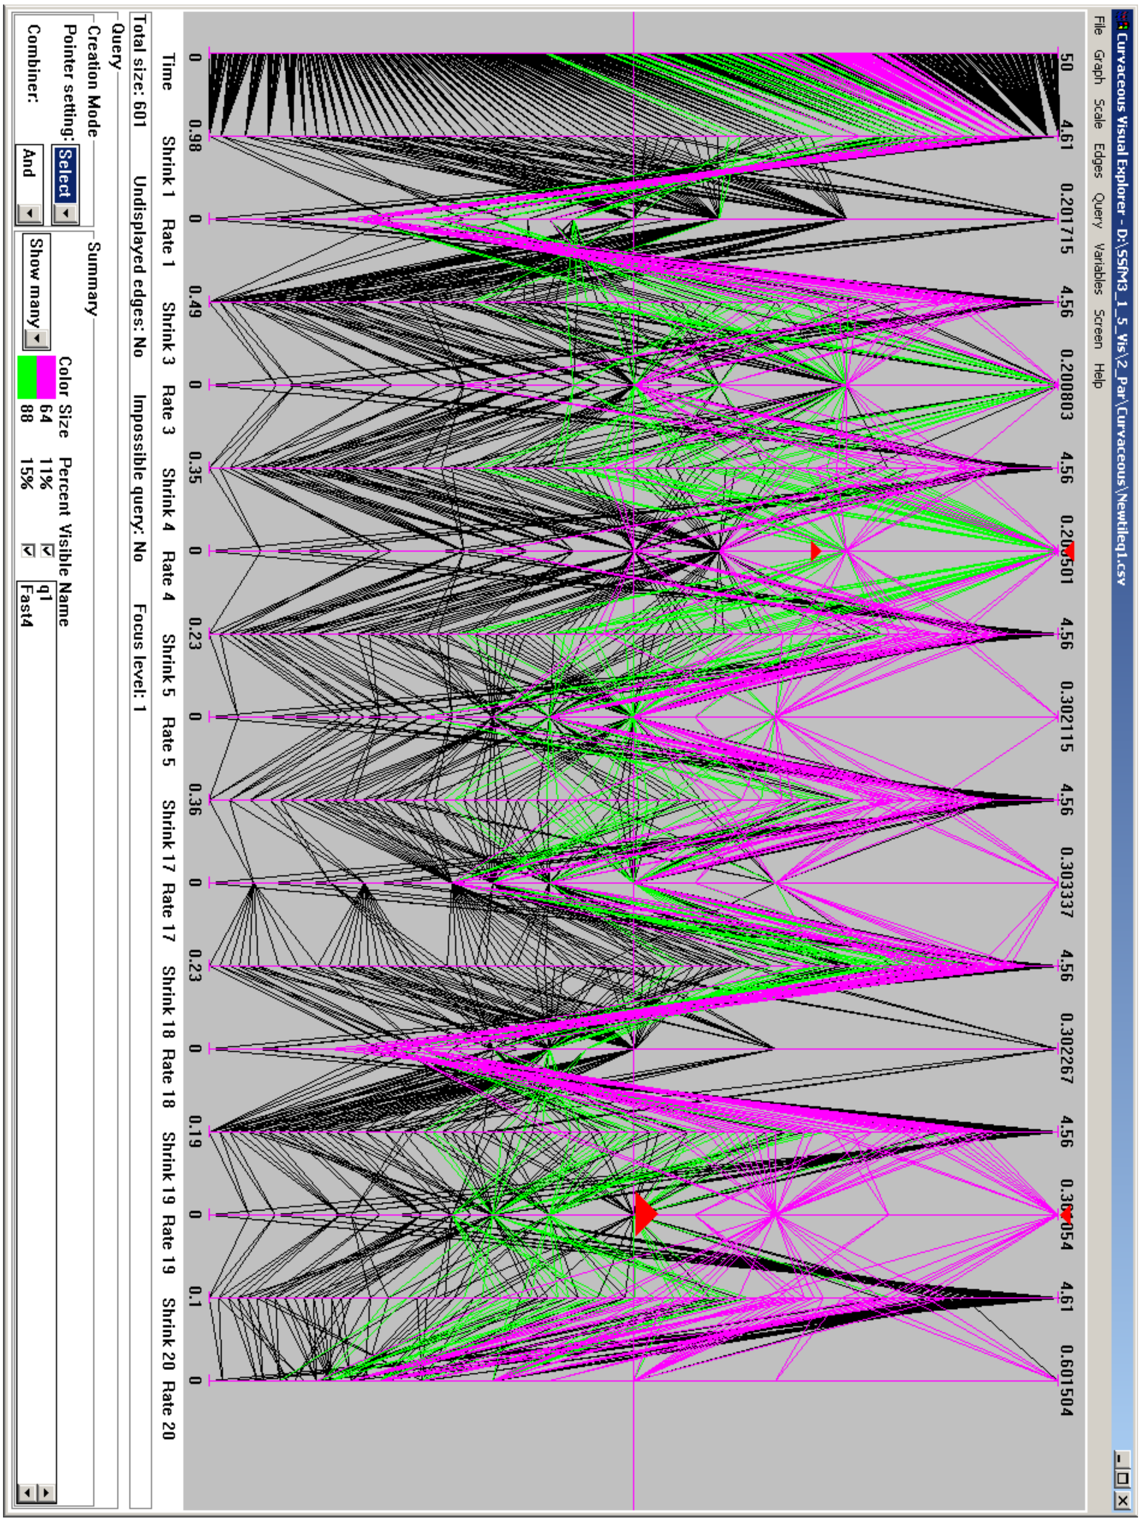

Figure 37: Data gathered after modifications to the dryer. Highlighted points are the fastest shrinkage rates for tiles 4 (green) and 19 (pink). Some of the fastest shrinkages occur at the same time.

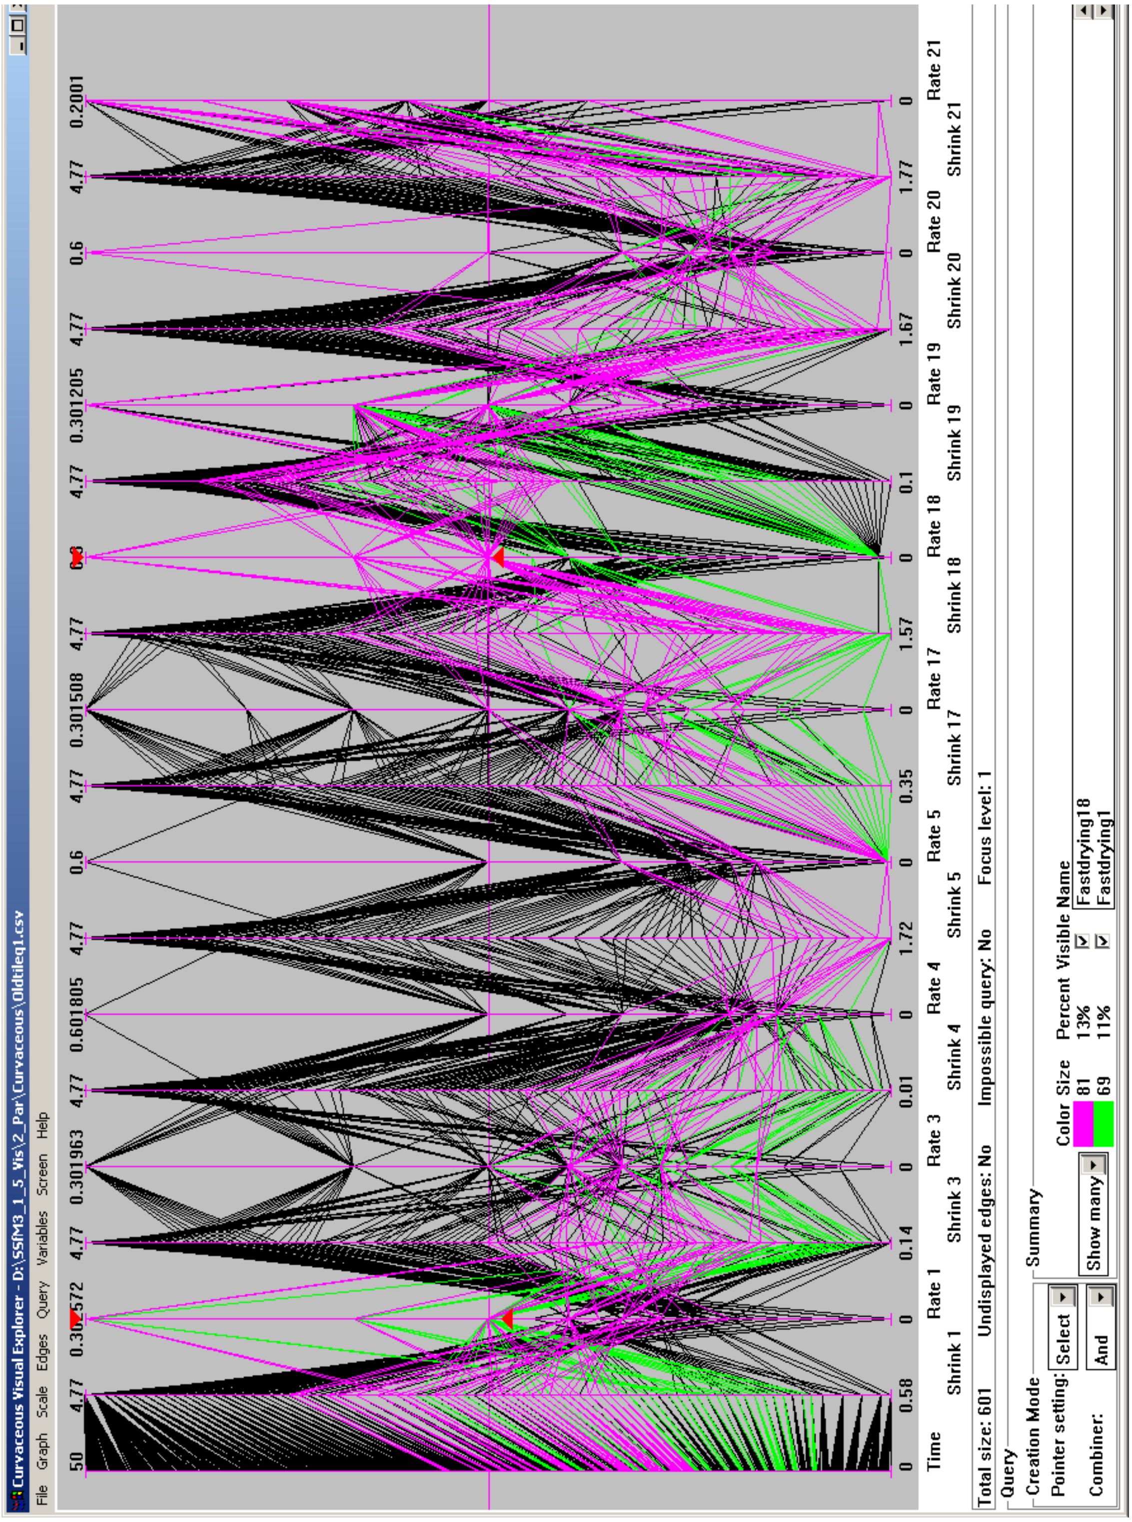

Figure 38: Data gathered before modifications to the dryer. Highlighted points are the fastest shrinkage rates for tiles 1 (green) and 18 (pink). The fastest shrinkages occur at different times.

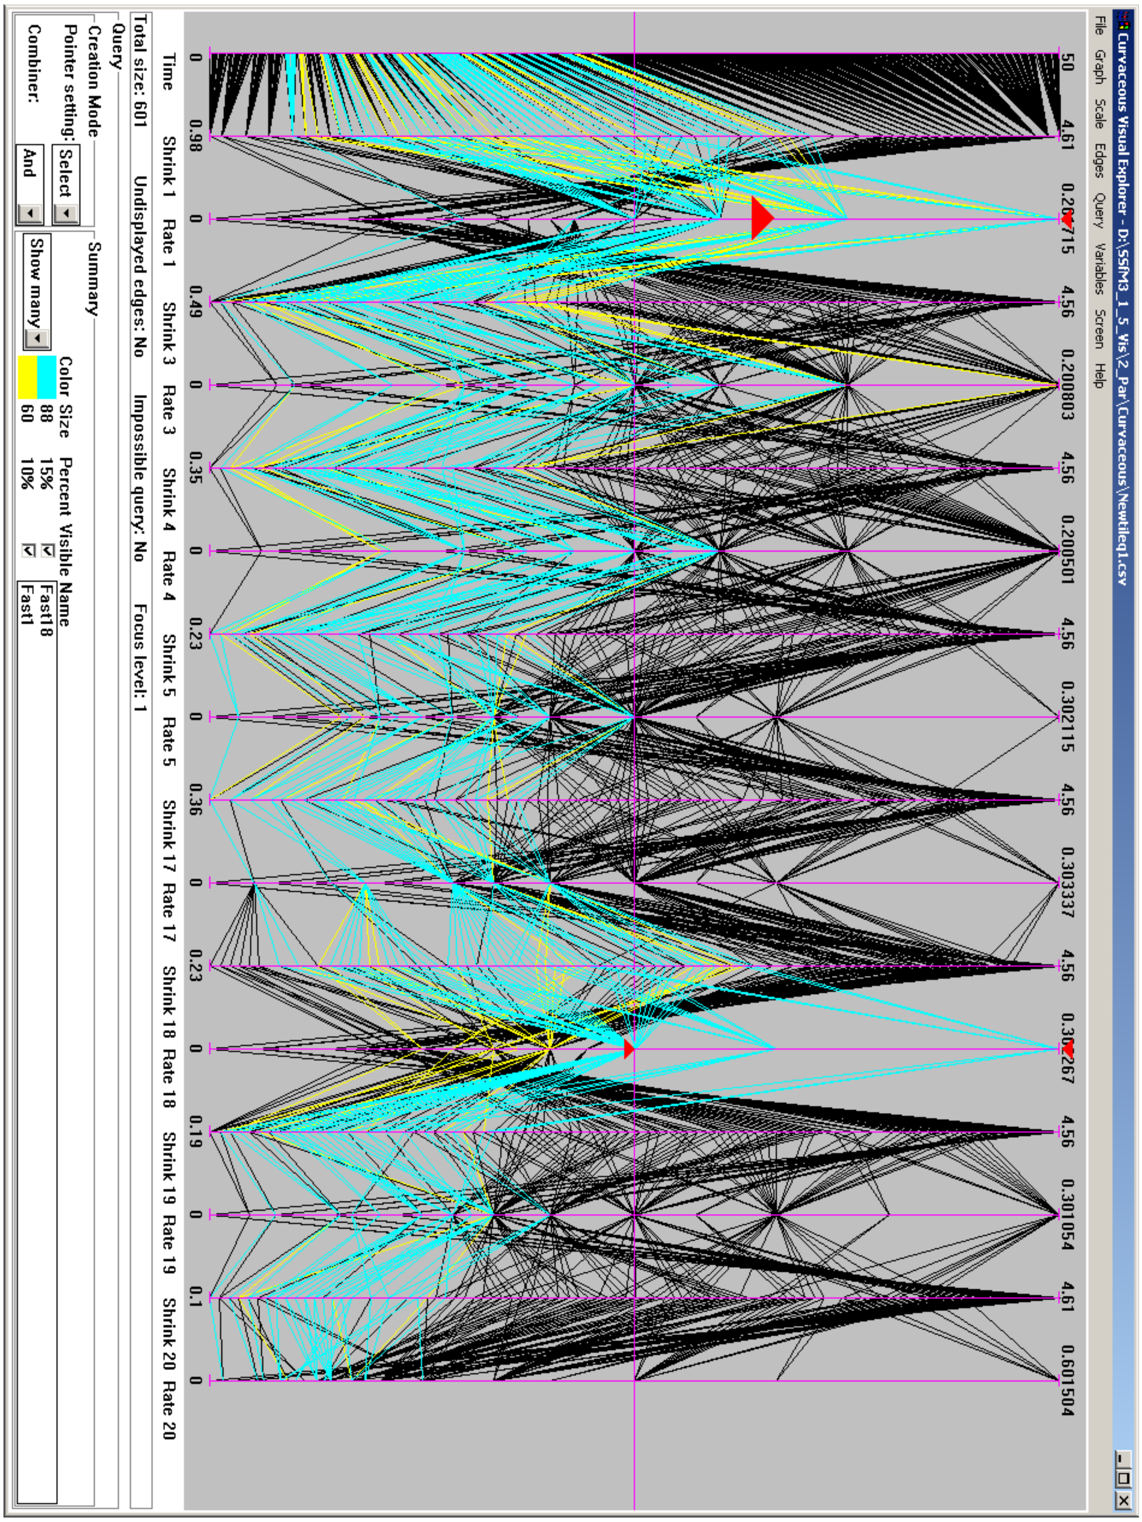

Figure 39: Data gathered after modifications to the dryer. Highlighted points are the fastest shrinkage rates for tiles 1 (yellow) and 18 (light blue). Some of the fastest shrinkages occur at the same time.

in yellow. The two sets of line segments are almost mutually exclusive, suggesting that the modifications to the dryer have not improved the top-to-bottom uniformity on that side of the dryer car.

Figure 42 shows similar results for the side of the car furthest from the fan before the modifications. The highlighted points are the fastest 17% of the shrinkage rates for tile 19. The fastest rates for tile 19 do not correspond to fast rates in any of the other tiles. In general, the fastest drying rates for tiles 18, 19, 20 and 21 do not occur simultaneously. This nonuniformity is probably because these tiles are further from the fan so the air flow over the tiles is less uniform than that of the other side of the dryer car.

Figure 43 shows results gathered in the same place after the modifications. The shrinkage sensor for tile 21 on this run did not produce any usable data, so only tiles 18, 19 and 20 can be compared. The modifications do not seem to have improved the uniformity. Tile 18 is still drying more quickly than tiles 19 and 20, and whilst the times of fastest shrinkage of tiles 19 and 20 overlap to some extent, they are no more similar than before the modifications.

Question 2: Effects of modifications

Some of the effects of the modifications have been shown in the figures of the previous section, particularly figures 36 and 37, which showed that the horizontal uniformity within the dryer car had improved. Figure 44 shows the full set of temperature measurements before and after the modifications on a single plot. The axis labels on this figure have three parts. Old means that the measurements were made before modification, and New means they were made after modifications. Tdry is an air temperature, Twet is a wet bulb temperature, Ttile is a tile temperature, and RH is a relative humidity measurement. The final number on each axis label identifies the sensor position (1 or 2, corresponding to tiles 4 and 19 in figure 35).

The axes have been ordered so that sets of measurements of the same quantity in the same position gathered before and after modification are adjacent, so for instance the axes labelled Old_RH_1 and New_RH_1 are adjacent. In each case, the “Old” axis comes before the “New” one. Each pair of axes that represent the same quantity at the same position has been scaled to run between the same limits. These choices make it easier to investigate the changes in ambient conditions caused by the modifications to the dryer.

In general, the line segments connecting two air temperatures measured and the same position or two tile temperatures measured at the same position are either horizontal or slope downwards (with a few exceptions). This feature means that in general the modifications have reduced the air and tile temperatures. The effect is particularly noticeable at the high temperature end of the scale. The line segments between these pairs of axes do not cross over one another very much, which means that the relationship between old temperature and new temperature is close to monotonic.

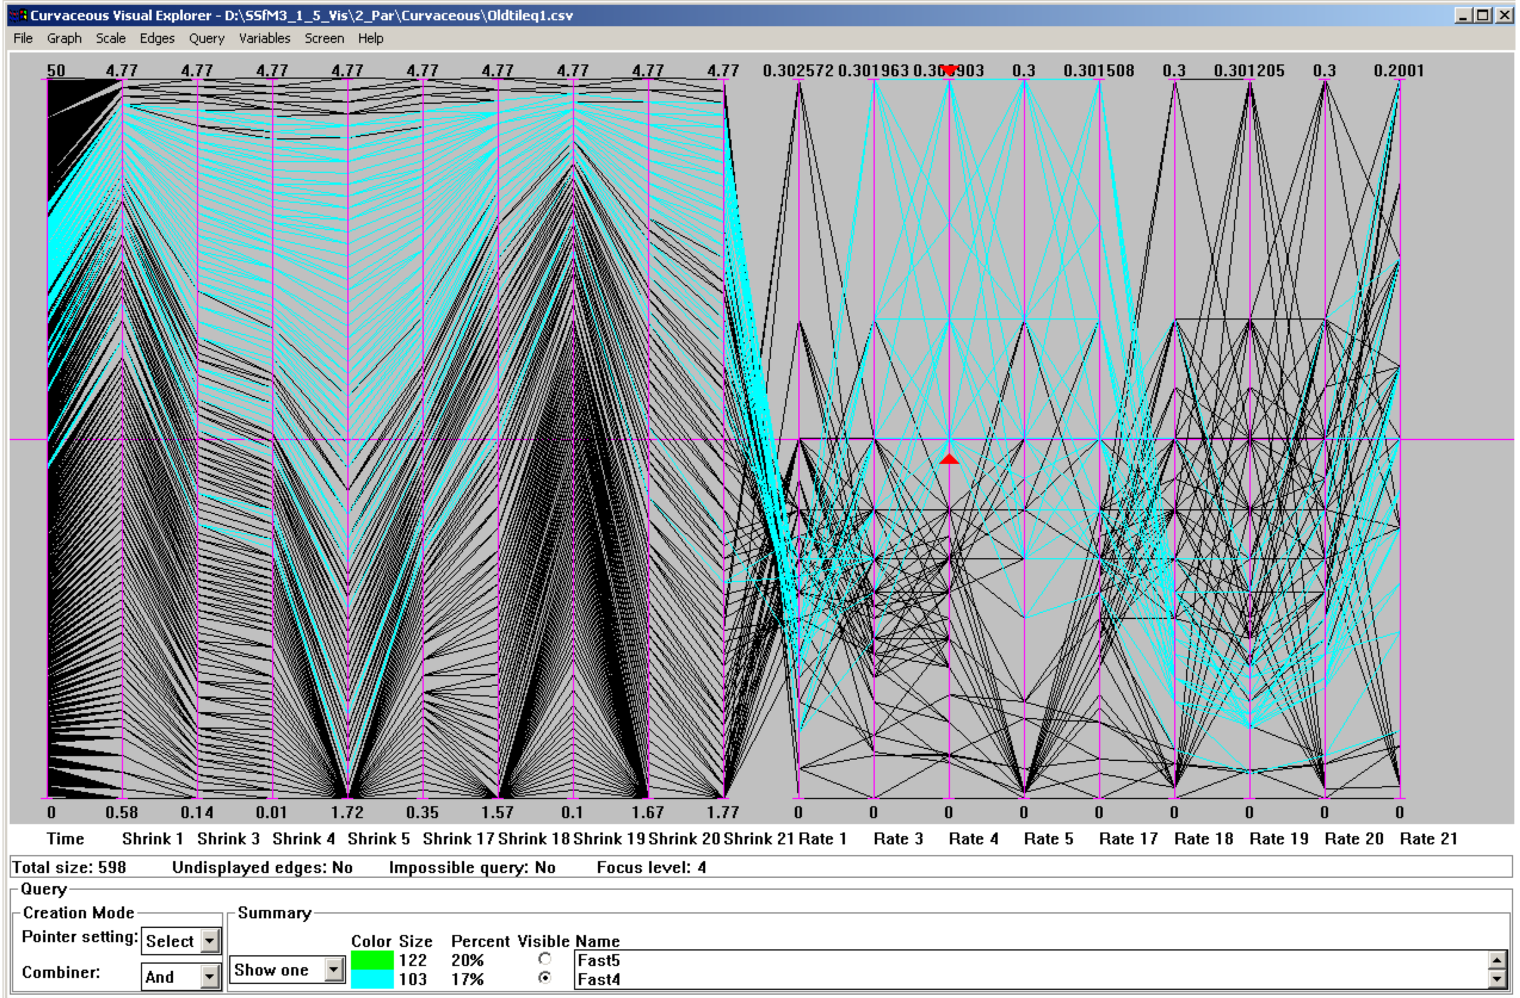

Figure 40: Data gathered before modifications to the dryer. Highlighted points are the fastest shrinkage rates for tile 4 (light blue). The uniformity of the side of the dryer car nearest to the fan is quite good, other than tile 1.

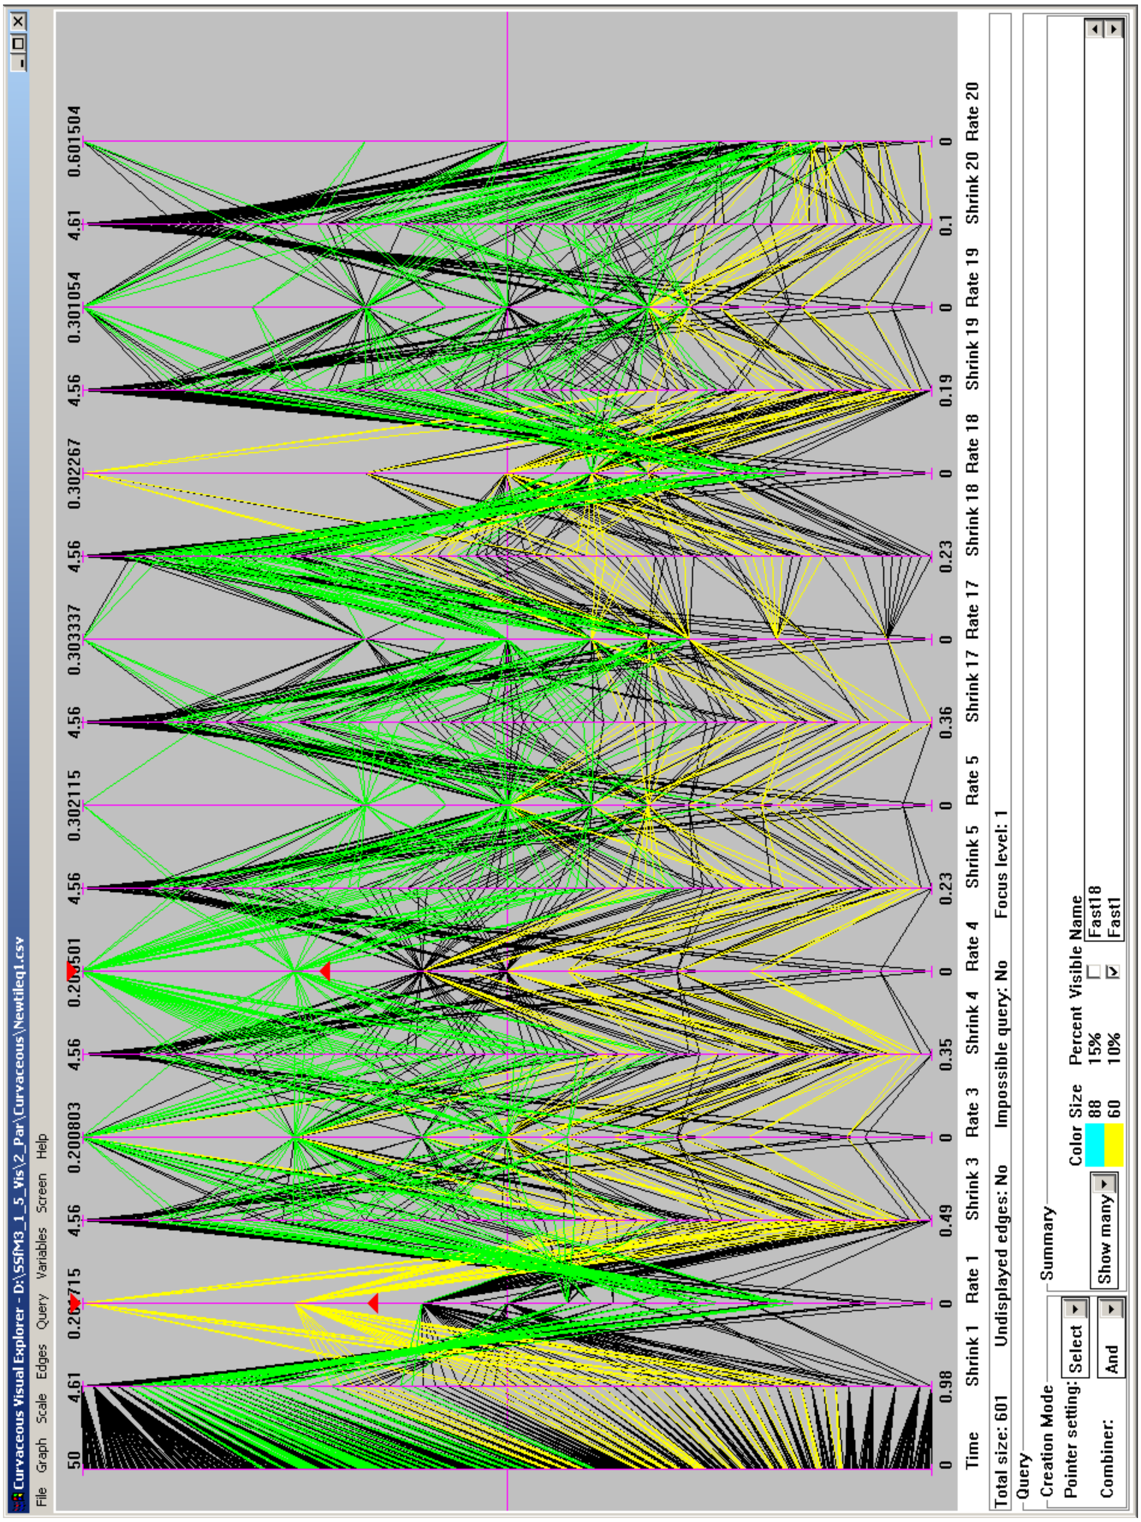

Figure 41: Data gathered after modifications to the dryer. Highlighted points are the fastest shrinkage rates for tiles 1 (yellow) and 4 (green). Even after modification, tile 1 dries faster than the other tiles on that side of the car.

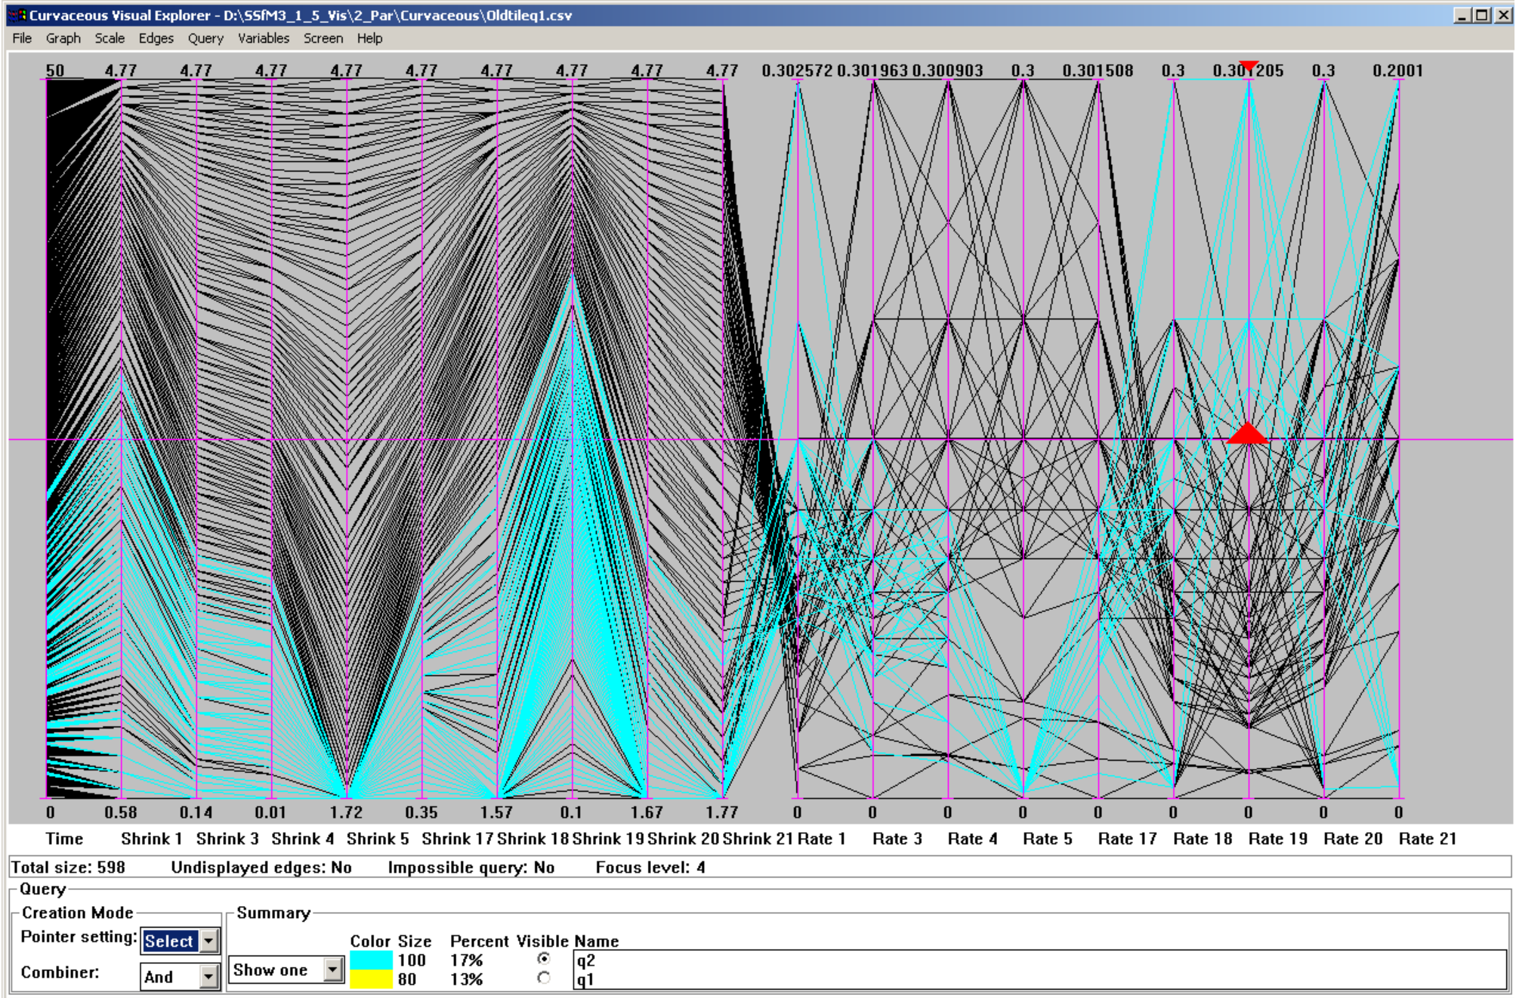

Figure 42: Data gathered before modifications to the dryer. Highlighted points are the fastest shrinkage rates for tile 19 (light blue). The uniformity of the side of the dryer car furthest from the fan is not very good.

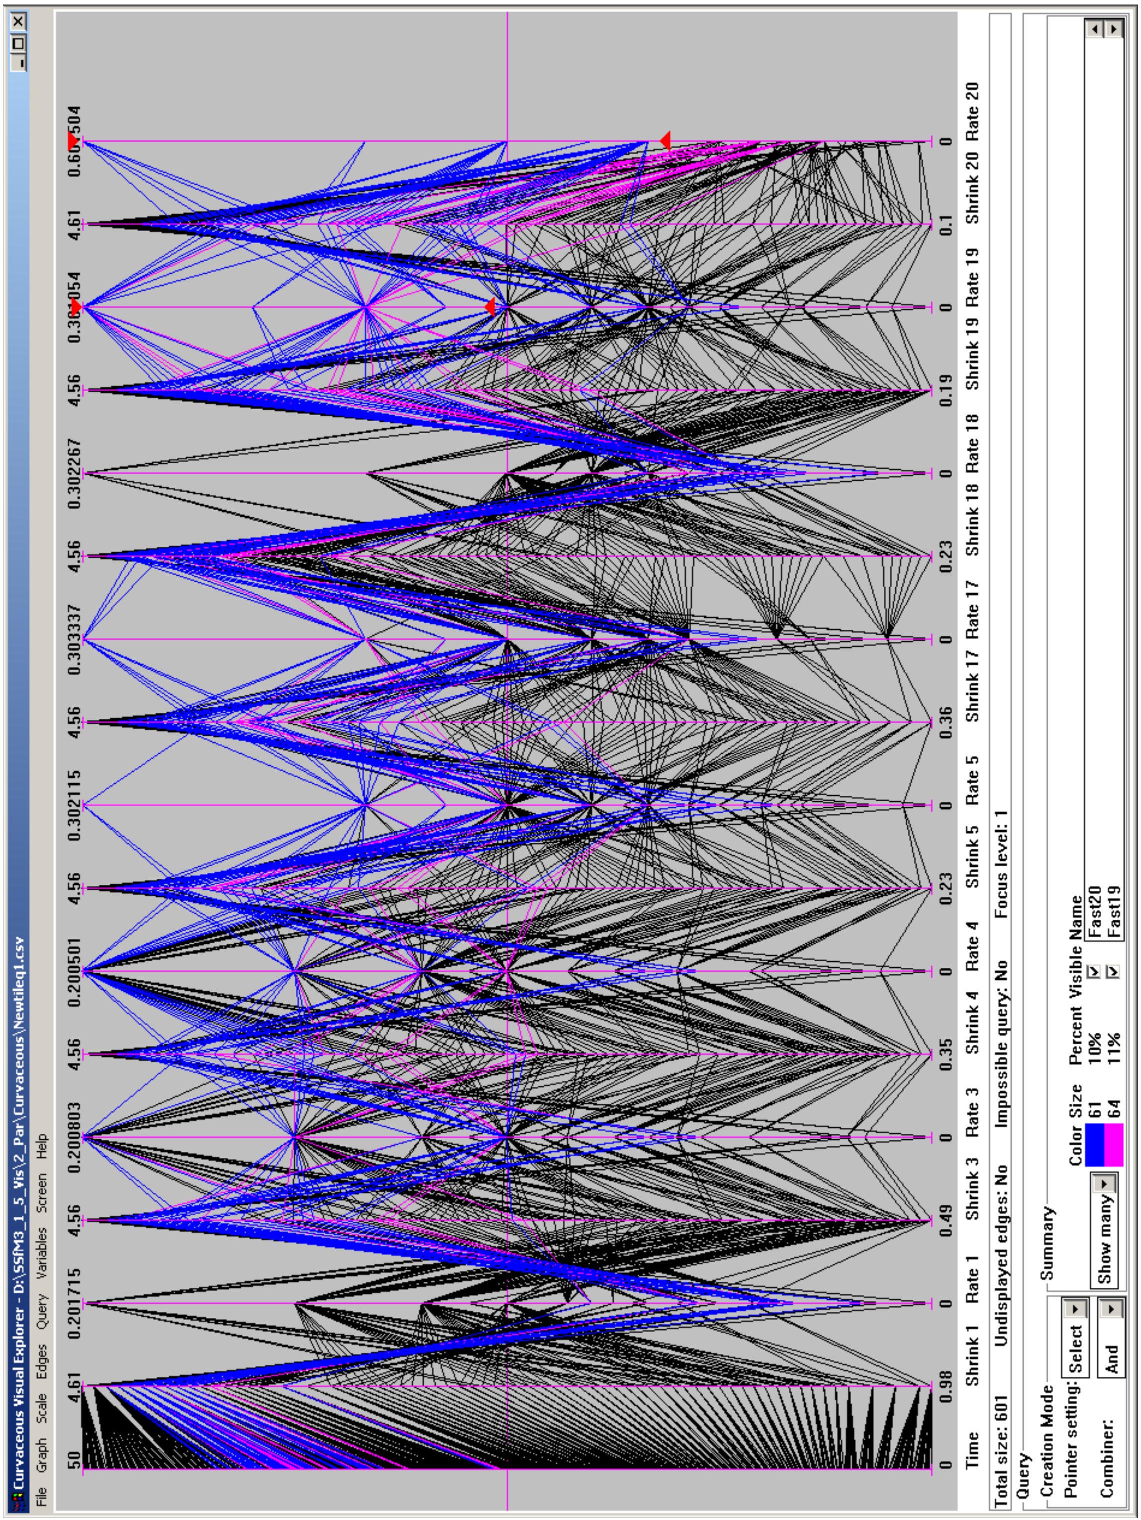

Figure 43: Data gathered after modifications to the dryer. Highlighted points are the fastest shrinkage rates for tiles 19 (pink) and 20 (dark blue). The modifications have not improved the vertical uniformity.

The line segments between pairs of wet bulb measurements cross over one another to a greater degree, particularly at the higher end of the scale. The highlighted points in figure 44 are the line segments that are causing most of the intersection for sensor 1, and the same effect is observed in the line segments corresponding to wet bulb sensor 2. The intersections occur at the end of the drying cycle (shown by the position on the time axis), when the air temperature, and the drop in air temperature noted above, is at its highest. If the air temperature had remained constant, the drop in wet bulb temperature would mean that the air was more dry, i.e. a lower relative humidity. However, since the air temperature has dropped by a significant amount, the relative humidity actually rises slightly and so the air is wetter. Since these effects occur during the final stages of the drying cycle, it is unlikely that they have a significant effect on the drying of the tiles, although in general reduced air temperature and increased humidity result in slower drying.

Figure 45 shows all of the shrinkage rate data gathered before and after modification, following the removal of a few outliers (since the shrinkage data is so blocky, some over-estimates of shrinkage rate occur). On the axis labels, S denotes that the data were gathered before modification, N denotes that they were gathered after the modifications, and the number indicates which tile position the data were gathered at. The axes have been ordered and scaled such that it is easy to compare“before” and “after” data.

The main points to note are that there is no clear pattern in the line segments between each pair of axes (even when some points are highlighted), and that every axis has a relatively small (generally less than 20) number of shrinkage rate values on it. The discretised nature of the data makes looking for patterns difficult. The discretisation of the rates is caused by the discretisation of the shrinkage sensors. In future this problem may be alleviated if 12-bit loggers are used to collect the data.

Question 3: Drying rates

Figures 46 to 50 show all the data gathered at positions 4 and 19. The axes in figures 46 to 50 represent, from left to right:

- whether the measurement was made before (1) or after (2) the modifications,
- whether sensor set 1 (tile 4) or 2 (tile 19) was used,
- time of measurement,
- air temperature,
- wet bulb temperature,
- tile temperature,
- relative humidity,

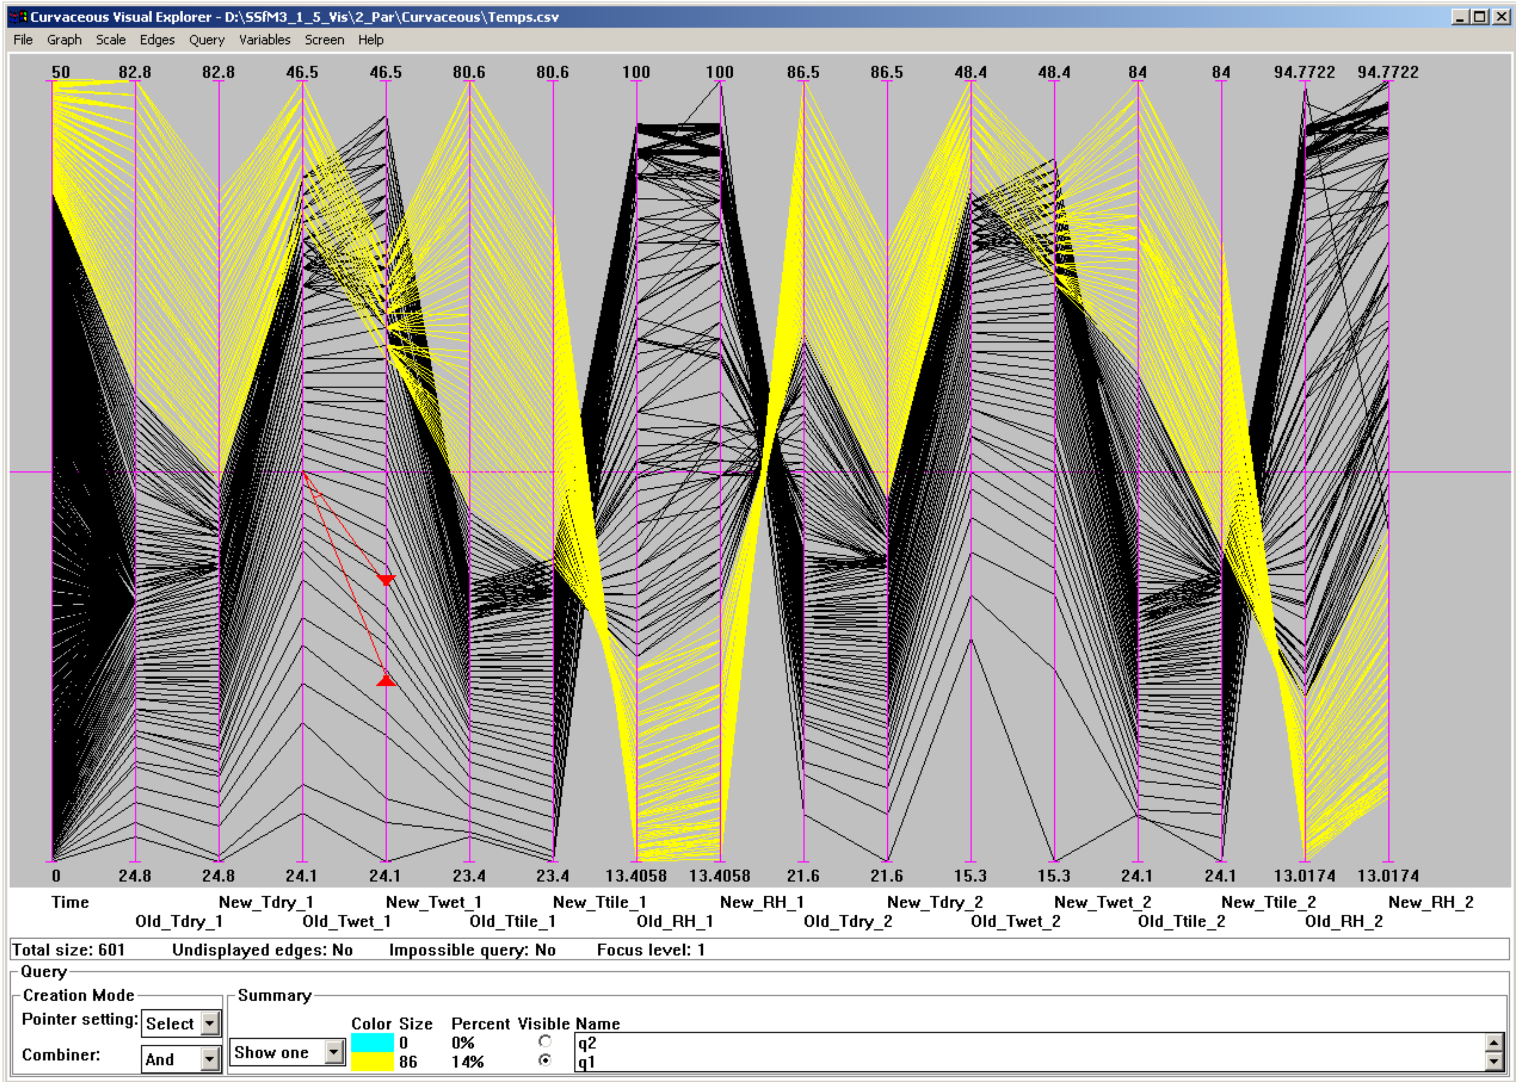

Figure 44: All the temperature and humidity data gathered before and after the modifications. The highlighted points show the main areas of intersection between line segments.

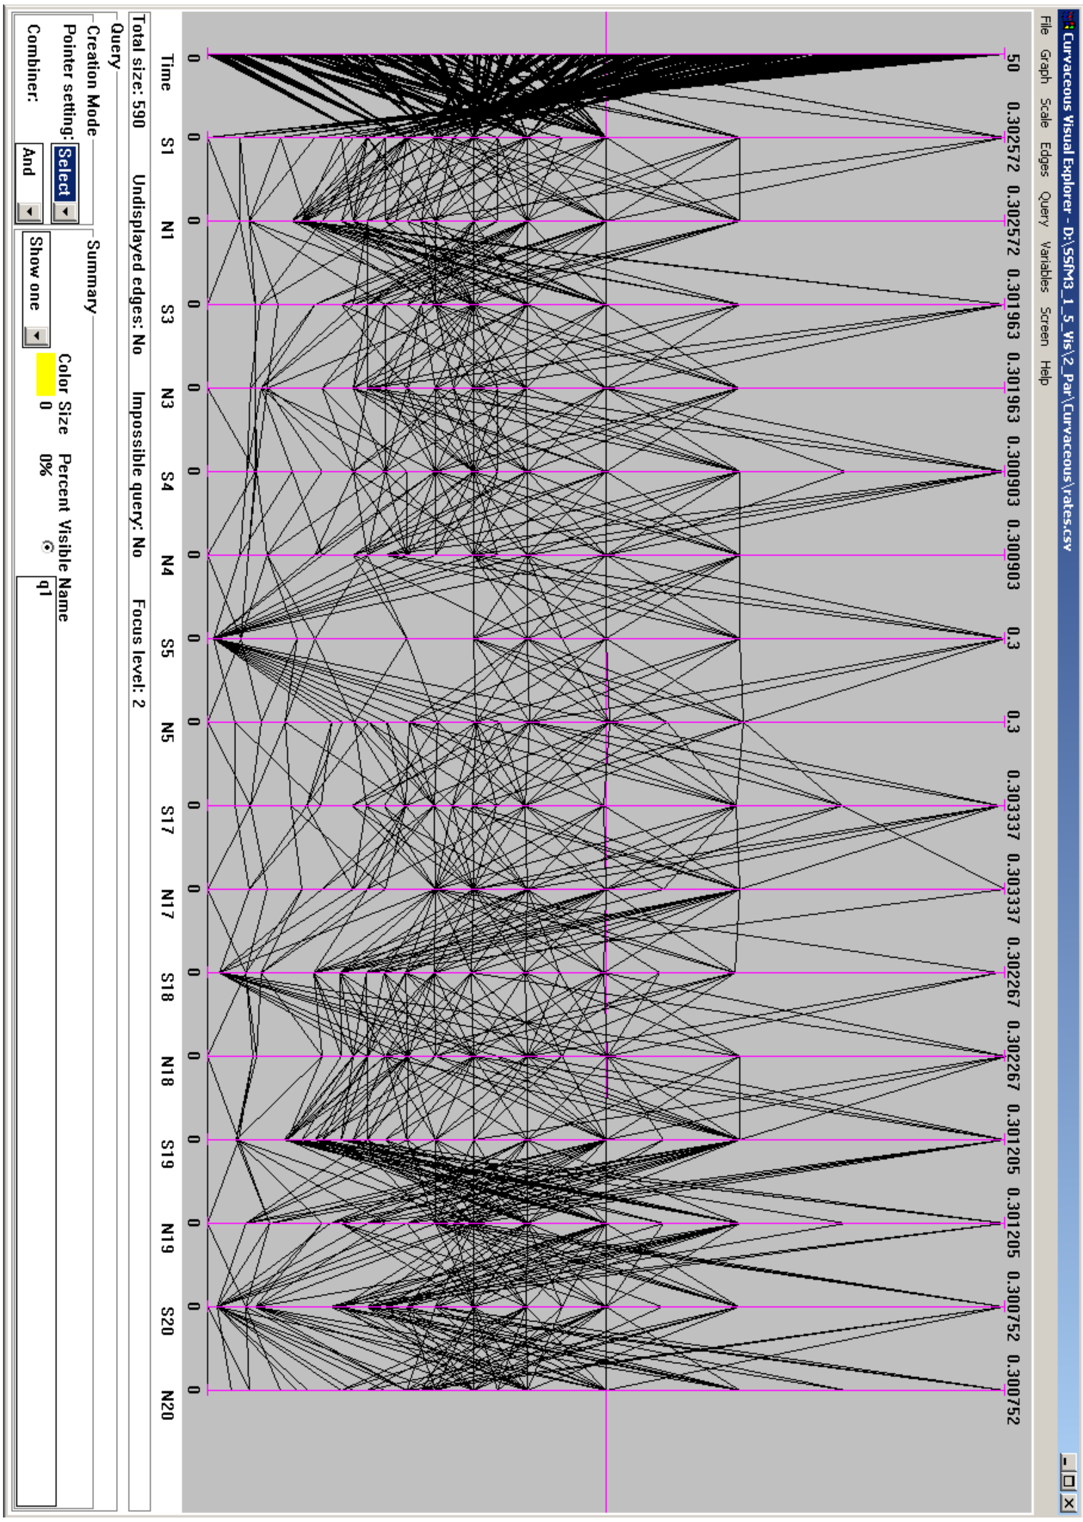

Figure 45: All the drying data gathered before and after the modifications.

- tile shrinkage, and
- shrinkage rate.

Note that the rate and wet bulb scales have been altered to remove two outliers and that the relative humidity scale runs in reverse order.

Figures 46 to 50 have been generated by splitting the drying time into five 10 hour segments and highlighting each one separately. The figures show that a range of shrinkage rates occur within each time band, and there is little or no correlation between any of the temperatures and the shrinkage rate. There may be too many points highlighted simultaneously for any clear relationship to show.

To investigate the relationship between temperature and shrinkage rate further, figure 51 highlights the slowest 25% of the drying rates (excluding zero). It should be noted that once the tile reached its final shrinkage limit at the end of drying, the drying rate was taken to be zero and so if zero was included, the plot would be biased towards the end of the drying cycle. Figure 52 highlights the fastest 21% of the drying rates in the same way.

These figures give very little information that could not be obtained from looking at a straightforward line plot of the four data sets shown. In particular, figures 17 and 18 do not give any surprising insights into drying rates. Interestingly, further subdivision of figure 18 indicates that of the fastest 21% of the drying rates, 13% occurred before the alterations and 8% occurred after, which implies that the alterations may have slowed the drying down.

5.2.4 Conclusions

The parallel coordinate technique has been used to visualise data gathered from a clay tile dryer on two separate occasions, before and after the dryer was modified. The visualisations have shown that the modifications have improved the drying rate uniformity of the tile on the side of the dryer car nearest to the fans, but that they have had little effect on the other side of the dryer car. The modifications generally reduced the air and tile temperatures for a given time and position.

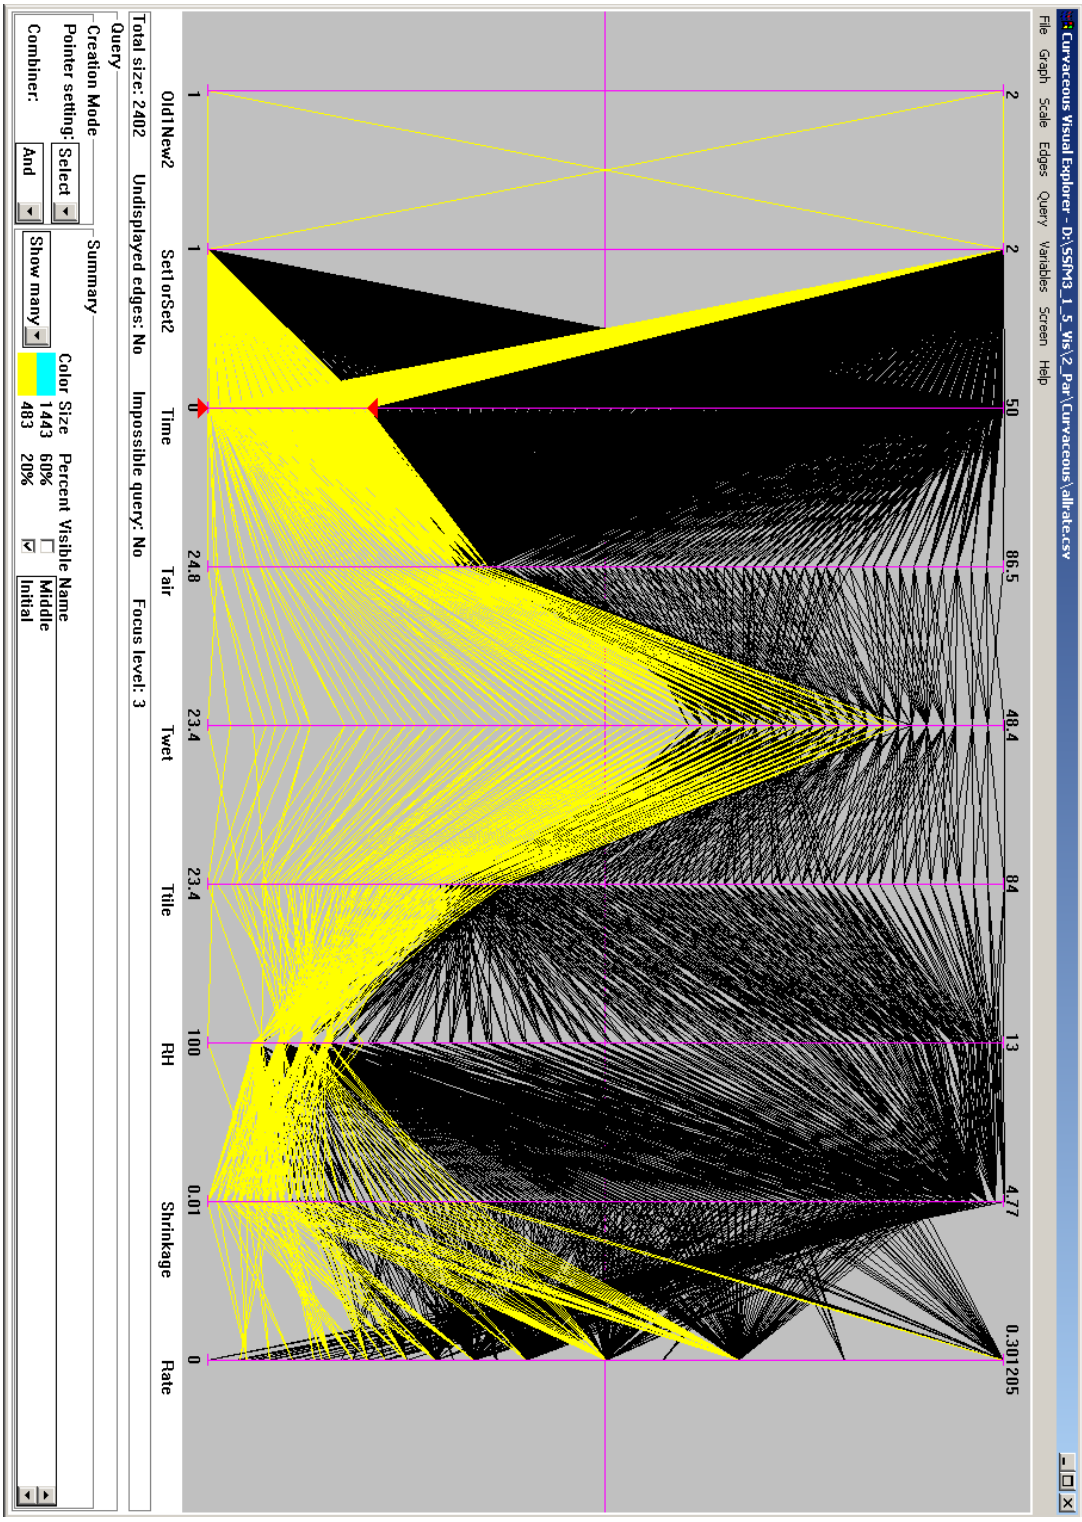

Figure 46: All ambient, shrinkage, and shrinkage rate data gathered before and after the modifications from tiles 4 and 19. Highlighted points are the first 10 hours of each drying cycle.

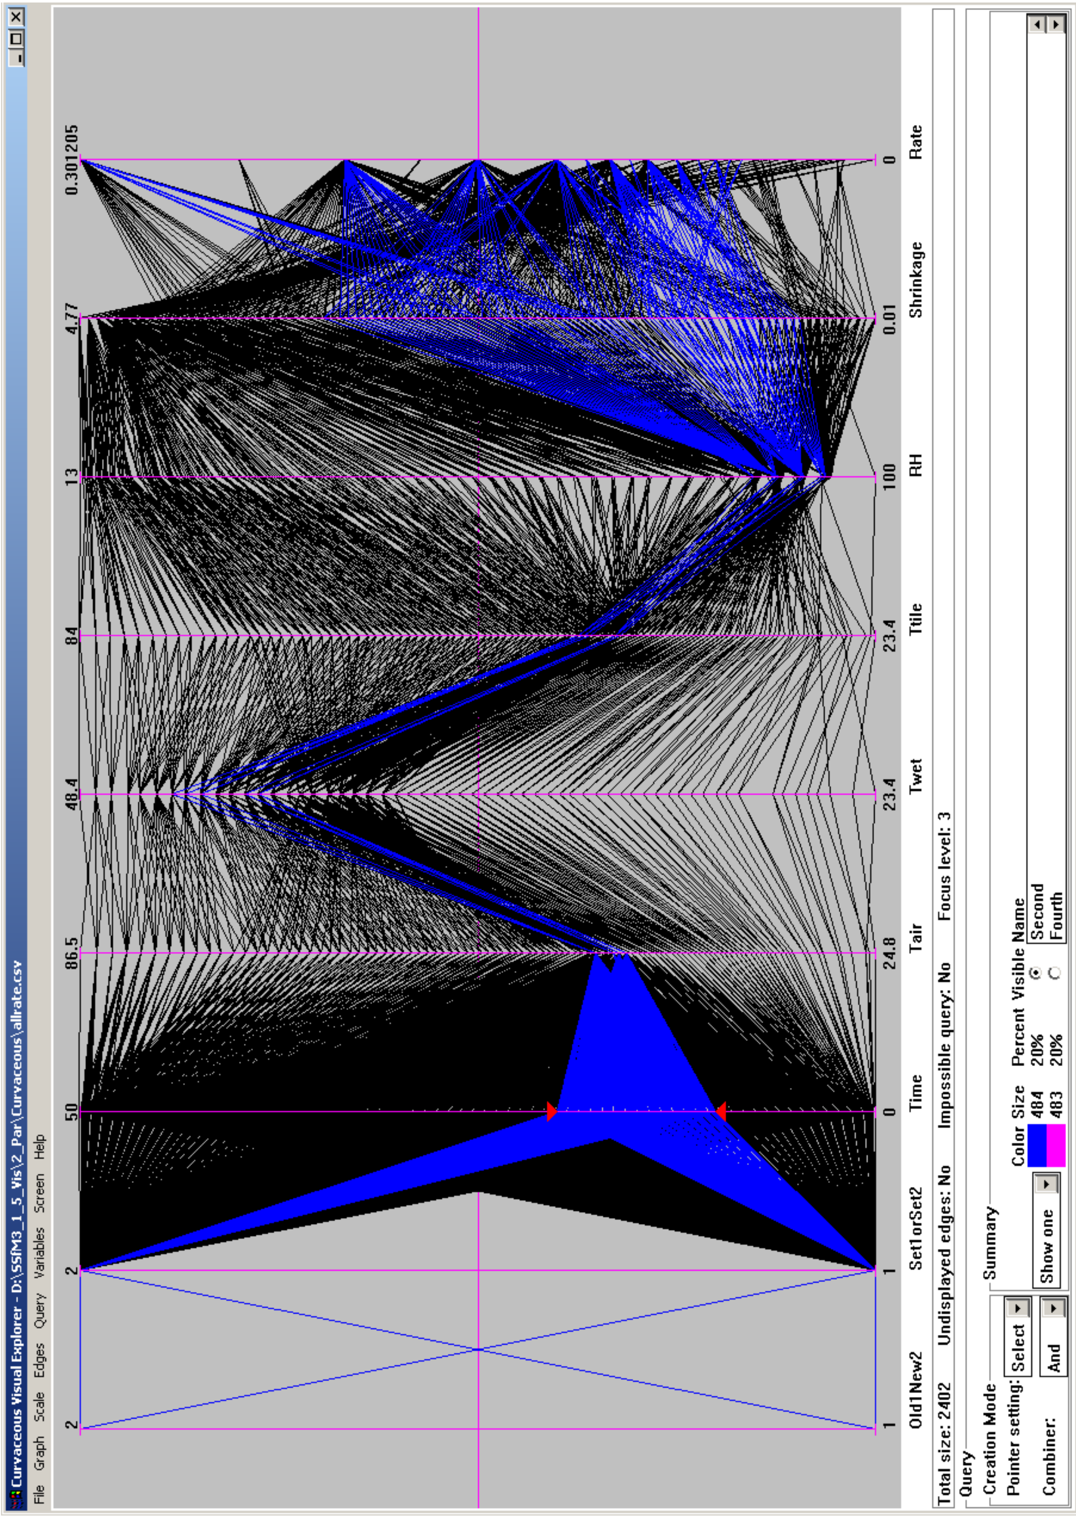

Figure 47: All ambient, shrinkage, and shrinkage rate data gathered before and after the modifications from tiles 4 and 19. Highlighted points are the second 10 hours of each drying cycle.

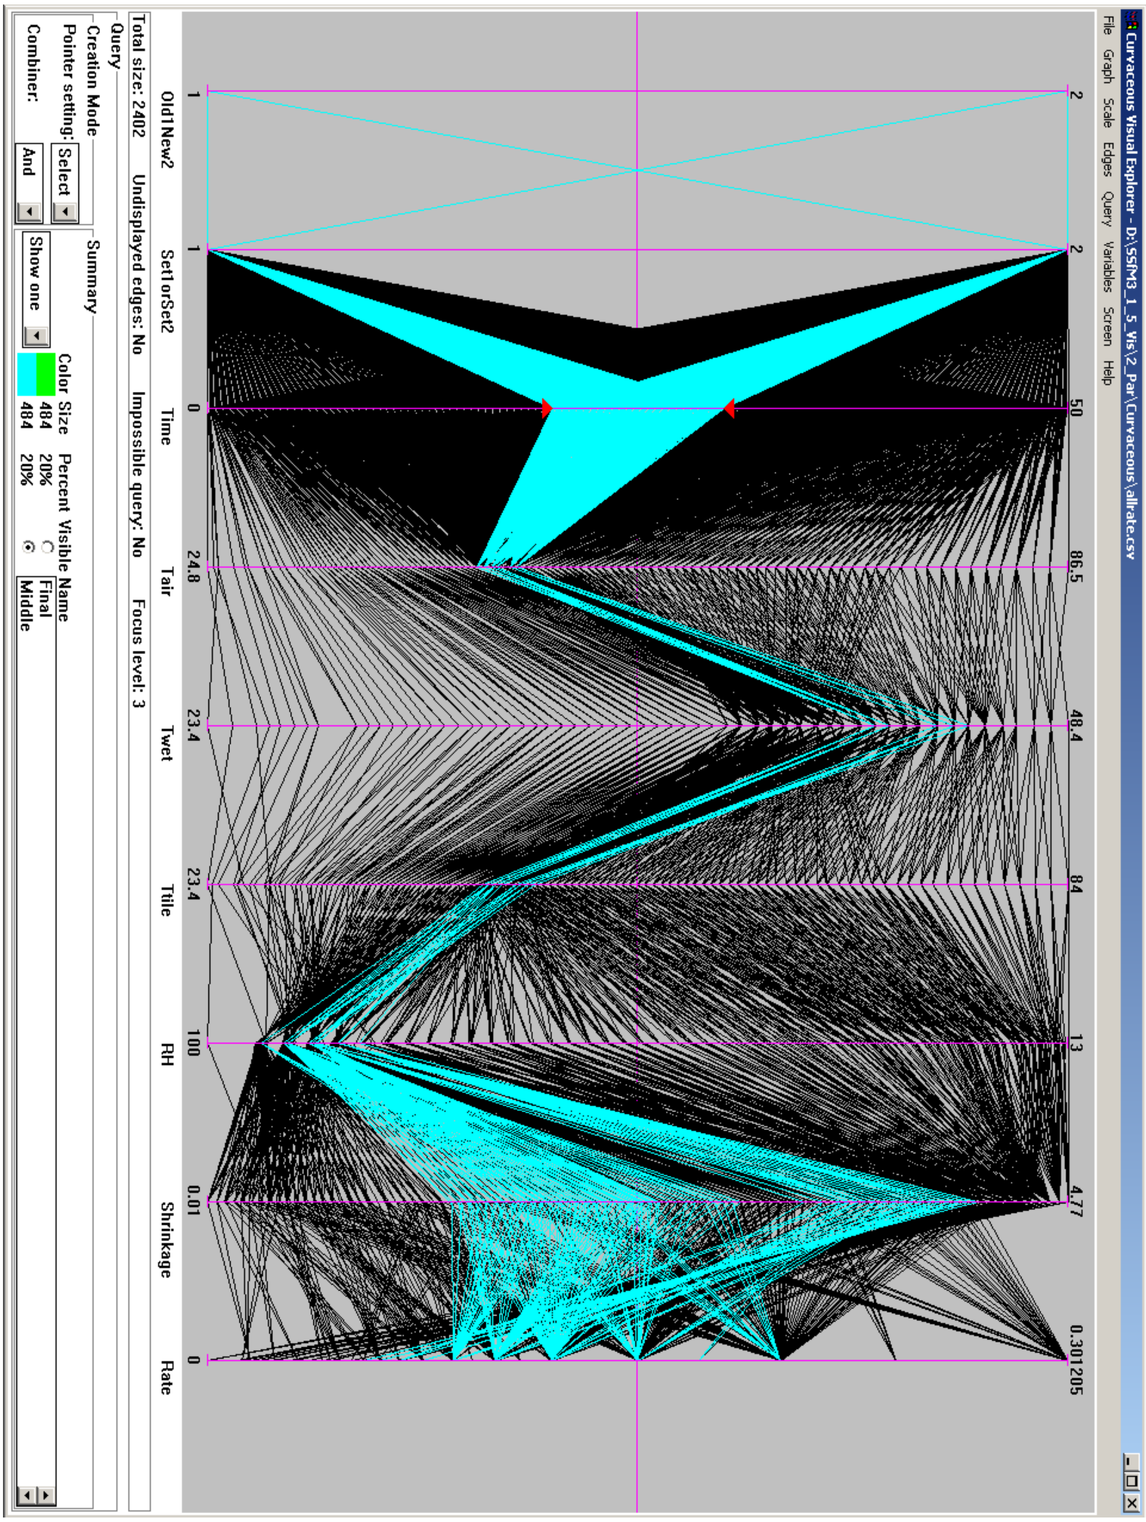

Figure 48: All ambient, shrinkage, and shrinkage rate data gathered before and after the modifications from tiles 4 and 19. Highlighted points are the third 10 hours of each drying cycle.

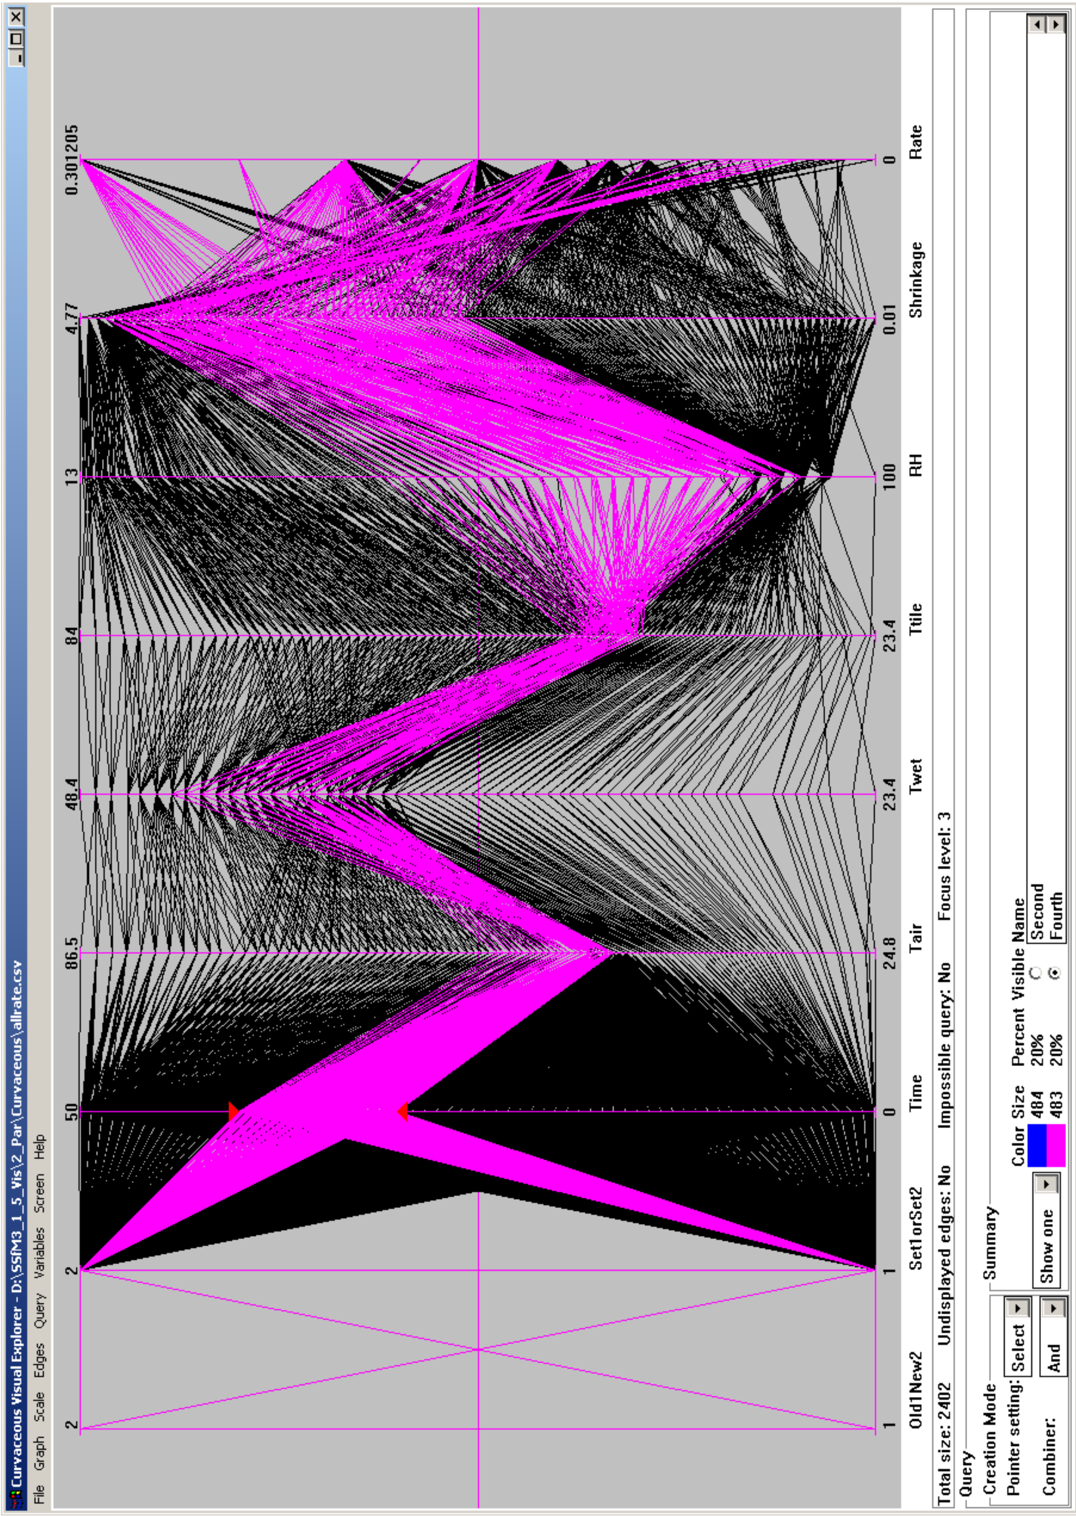

Figure 49: All ambient, shrinkage, and shrinkage rate data gathered before and after the modifications from tiles 4 and 19. Highlighted points are the fourth 10 hours of each drying cycle.

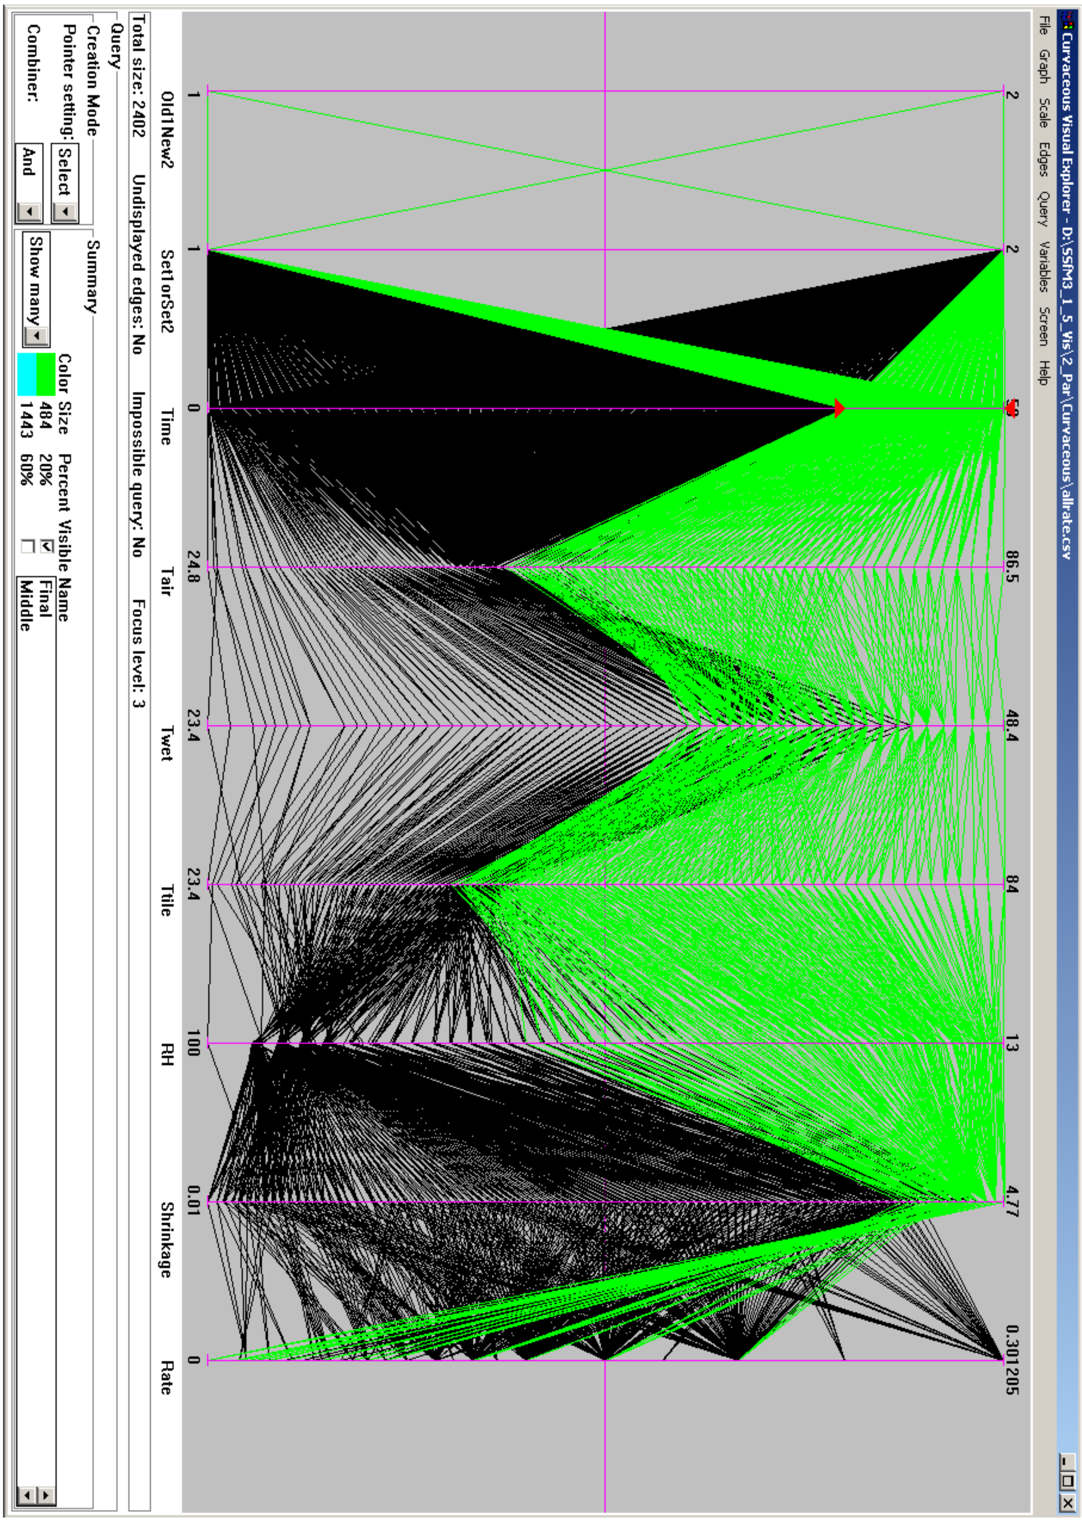

Figure 50: All ambient, shrinkage, and shrinkage rate data gathered before and after the modifications from tiles 4 and 19. Highlighted points are the final 10 hours of each drying cycle.

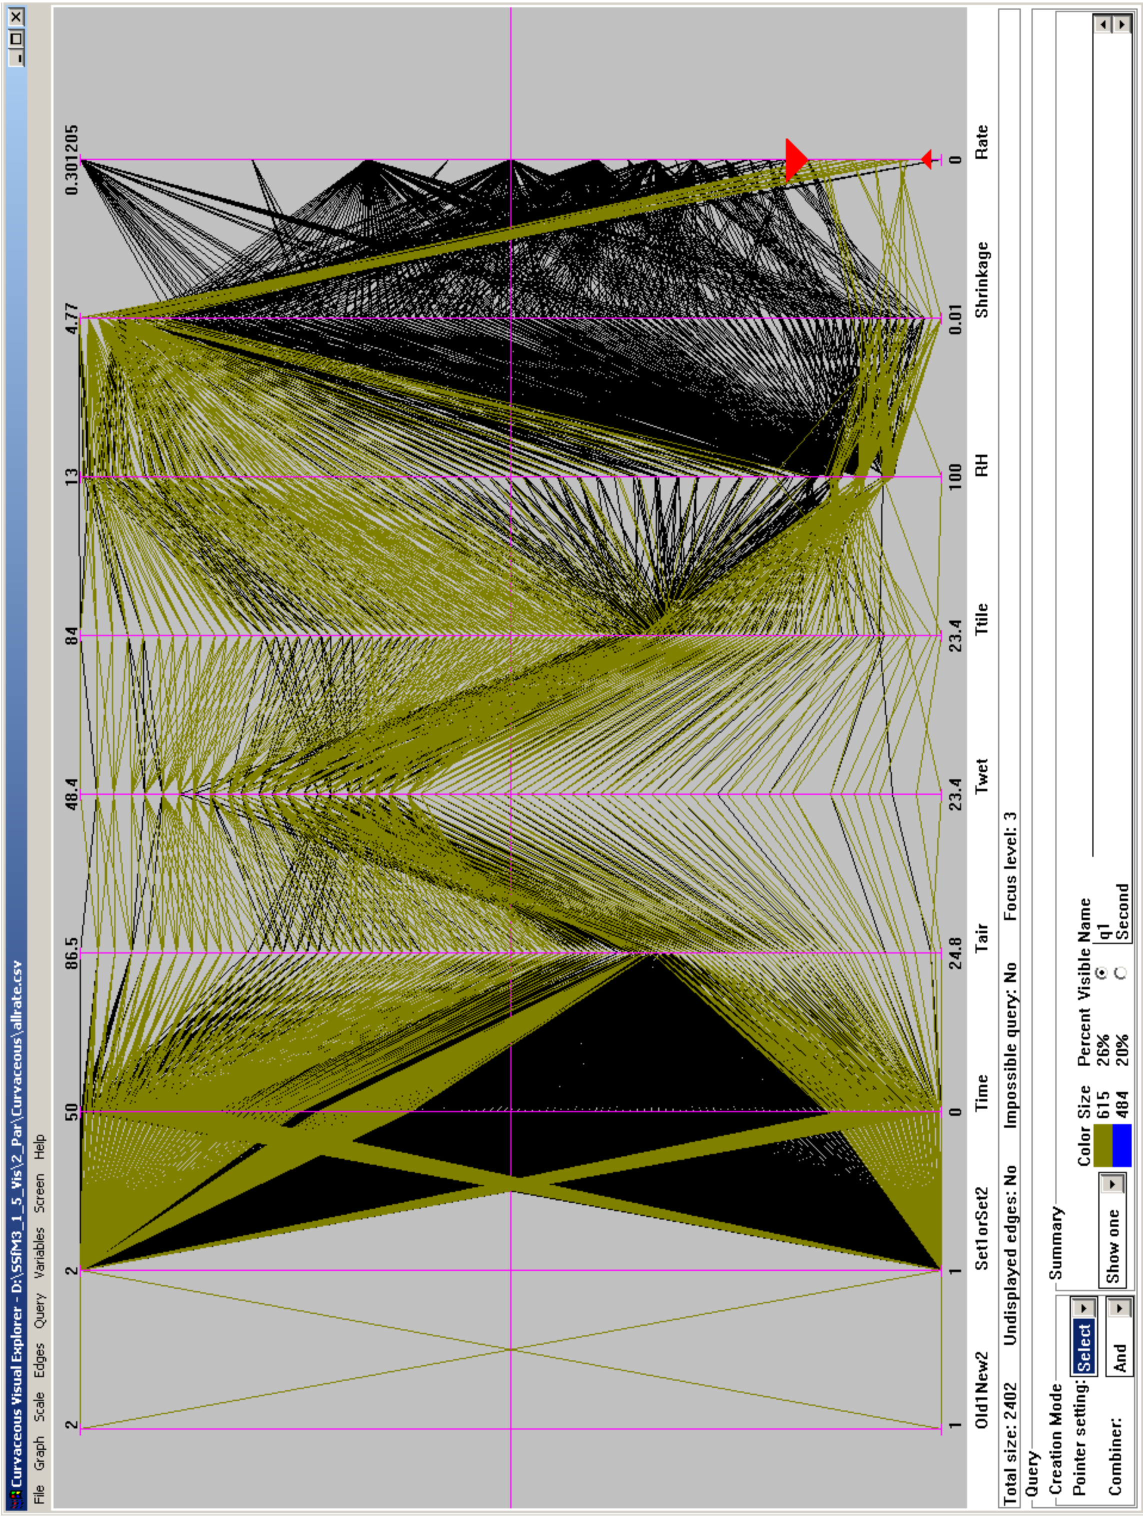

Figure 51: All ambient, shrinkage, and shrinkage rate data gathered before and after the modifications from tiles 4 and 19. Highlighted points are the slowest 25% of the drying rates.

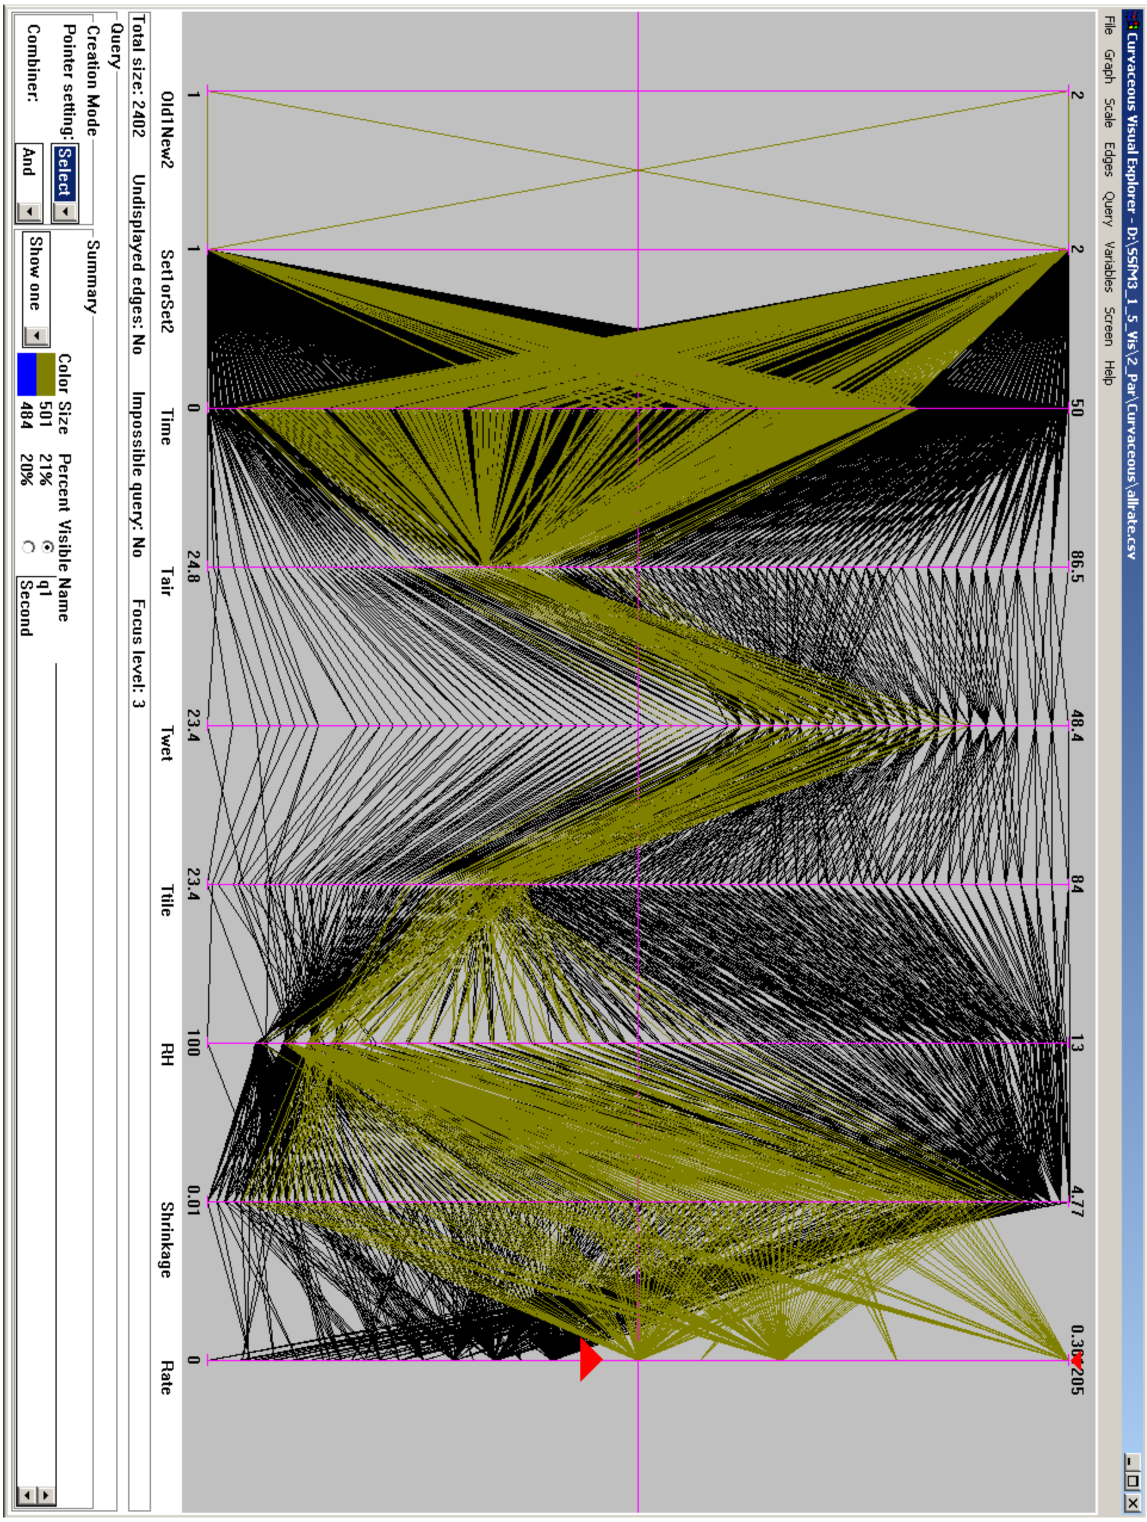

Figure 52: All ambient, shrinkage, and shrinkage rate data gathered before and after the modifications from tiles 4 and 19. Highlighted points are the fastest 21% of the drying rates.

6 Further reading and resources

The majority of this report was created by experimenting with software tools on data sets rather than by reading references. This approach was taken because much of the literature on parallel coordinates describes geometrical results that are of little relevance to the majority of problems in metrology. Hence this section does not provide a bibliography or reference list. Instead, we list some web resources that have been useful (some of which include more extensive reading lists), and provide website references for some of the tools described in the text.

As was mentioned in the introduction, parallel coordinates were developed in the 1980s by Prof. Alfred Inselberg. Prof. Inselberg has a website explaining the background to the parallel coordinate technique and listing some of his key publications on the topic. The website also includes a short tutorial on parallel coordinates, largely approached from the point of view of a geometrist. The website is at

http://www.math.tau.ac.il/%7Eaiisreal/.

Another web page, with a helpful introduction to the key concepts and the representation of common geometric shapes in multiple dimensions, is at

http://catt.bus.okstate.edu/jones98/parallel.html.

Websites associated with some of the software packages mentioned in this report are:

- CVE: http://www.curvaceous.com
- Parvis: http://home.subnet.at/flo/mv/parvis/
- Xmdv: http://davis.wpi.edu/xmdv/vis_parcoord.html
- XGobi: http://www.public.iastate.edu/%7Edicook/xgobi/xgobi.html
- Orange Canvas: http://www.ailab.si/orange/download.asp
- CrystalVision: ftp://www.galaxy.gmu.edu/pub/software/

Some papers that have proved of interest during this report are:

- *Parallel coordinates: survey of recent results, A. Inselberg, B. Dimsdale, A. Chatterjee, and C.-K. Hung, Proceedings of SPIE – Volume 1913 Human Vision, Visual Processing, and Digital Display IV, J. P. Allebach, B. E. Rogowitz (Eds), pp. 582-599, 1993.*
- *On Some Mathematics for Visualizing High Dimensional Data, E. J. Wegman and J. L. Solka, Sankhya: the Indian Journal of Statistics, v 64, Series A, Pt 2, pp429-452, 2002.*

- *Geometry unifies process control, production control and alarm management, R. Brooks, J. Wilson and R. Thorpe, Computing & Control Engineering Magazine, February 2004.*

Finally, the work described in this report has suggested several other areas of metrology that may benefit from visualisation using parallel coordinates. The main three such areas are

- Visualisation of multi-dimensional confidence regions
- Visualisation of the results of sensitivity analyses in multiple dimensions
- Visualisation of the output of wireless sensor networks

It is hoped that the work described here will benefit these areas in the future.

7 Appendix A: Common features of measurement data

This appendix will focus solely on pairwise comparisons of adjacent axes as a means of identifying features of measurement data. It will not consider features of multi-dimensional surfaces. It is important to note that whilst these features of measurement data can be helpful in identifying familiar concepts, the main advantage of the parallel coordinate technique is its ability to identify multi-dimensional relationships.

Suppose that a parallel coordinate plot has two axes displaying points $x_i, i = 1, 2, \ldots, N$, and $y_i, i = 1, 2, \ldots, N$, reading from left to right. Then the ratio $\frac{y_i - y_{i-1}}{x_i - x_{i-1}}$ is an approximation to the gradient of the function $y(x)$. This feature means that if the x_i are evenly spaced, then the user can get an idea of how the gradient is changing by looking at the way $y_i - y_{i-1}$ is varying with i.

Some common features of measurement data will be examined using seven functions defined at the points $x_k, k = 0, 1, \ldots, 100$ where $x_k = \frac{k}{100}$. The functions have been chosen to illustrate a particular feature that can be seen in a parallel coordinate plots. The seven features are linearity, monotonic nonlinearity, minimum points, inflexion points, multiple minima/maxima, convergence to a single value, and convergence to a line. The functions used for each feature will be explained during the example.

It should be stressed that these examples are easy to interpret because they are smooth functions sampled at small equal intervals. Real measurement data will not be so straightforward to interpret, but will have approximately the same features as some of those shown here. The data sets used in the case studies illustrate the differences between smooth continuous analytic functions and noisy digitised measurement data.

7.1 Linearity

Linearity is one of the easiest relationships to spot using a parallel coordinate plot. As was stated above, if the x_i are evenly spaced, then the variation in $y_i - y_{i-1}$ is a reflection of the gradient of a normal X-Y plot of the two variables being examined. If the two variables are linearly related, their X-Y plot would fit a straight line of constant gradient. Hence if two variables are linearly related, the y_i will be evenly spaced, and hence the line segments linking the x_i and the y_i will be parallel.

The function used to illustrate linearity here is $y = 2x - 1$, which is plotted as an X-Y plot in Figure 53. The equivalent parallel coordinate plot is shown in Figure 54(a). All of the line segments are parallel, as was expected. The line segments are horizontal in this example, because the function is exactly linear over the entire range of y values, but parallel line segments at an angle to the horizontal also represent a linear relationship between the variables. The axes can always be translated to make the line segments horizontal and parallel. If the data values have a linear relationship plus a small amount of noise, there will be a few crossings of the line segments between the axes, but the line segments will be close to parallel.

An alternative way of looking at a linear relationship is shown in Figure 54(b). In this figure, the second axis has been reversed so that it runs from its maximum value to its minimum. The line segments now all cross at a single point. The two forms make it easy to spot likely linear and nearly linear relationships between variables whether the correlation is positive or negative.

The parallel line segments mean that any two axes that are linearly related are equivalent in as much as a plot that has y and x axes adjacent will look the same as one that has y and $ax + b$ axes adjacent. This is illustrated in Figure 56 using a monotonic function.

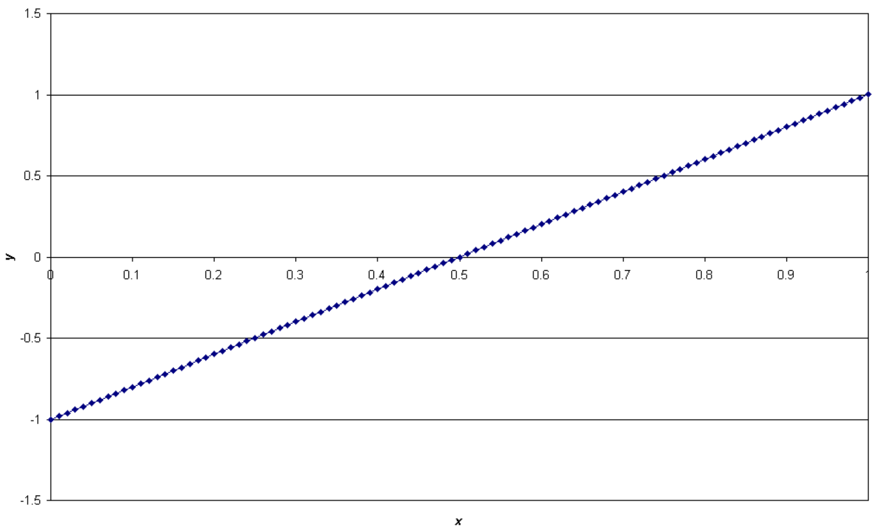

Figure 53: A linear function plotted as an X-Y plot.

7.2 Monotonic behaviour

Monotonic functions are also easy to identify in parallel coordinate plots because the line segments between the axes do not cross. The derivatives of monotonic functions are either always positive (increasing function) or negative (decreasing function), so $y_i - y_{i-1} \geq 0, i = 1, 2, \ldots, N$ for a monotonically increasing function or $y_i - y_{i-1} \leq 0, i = 1, 2, \ldots,$ N, for a decreasing one. Hence the spacing between consecutive y values will either be constantly non-negative or constantly non-positive.

The function used to illustrate this is $y = x^2$, plotted as an X-Y plot in Figure 55. The equivalent parallel coordinate plot is shown in Figure 56(a). Note that the gap between consecutive y values is increasing. Figure 56(b) shows three axes: x, x^2, and $2x - 1$ (from left to right). This plot illustrates that plotting x against x^2 produces the same results as

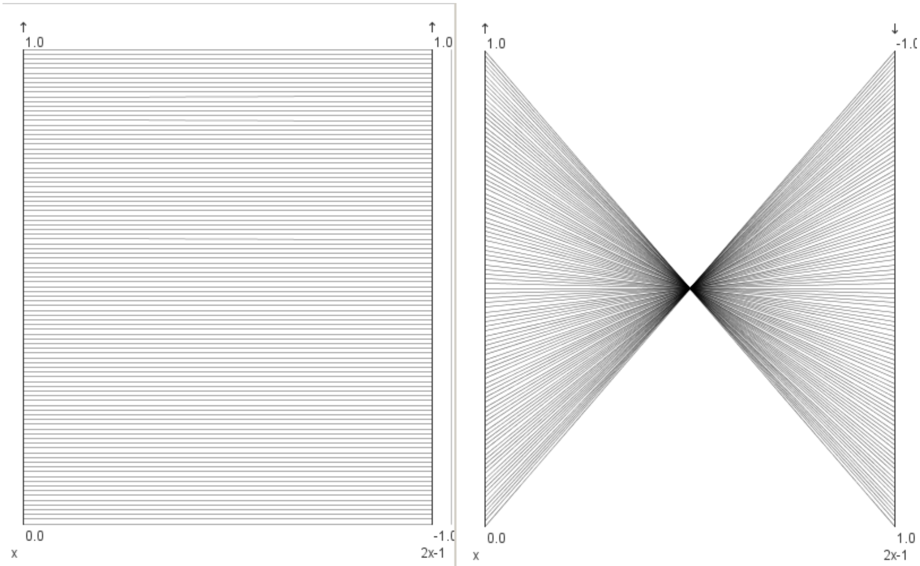

(a) (b)

Figure 54: The linear function shown in Figure 53, plotted as a parallel coordinate plot (a), and with one axis reversed (b).

plotting $2x - 1$ against x^2, as was stated in section 7.1.

7.3 Turning points

Maxima, minima, and points of inflection involve crossing of line segments in parallel coordinate plots, and hence are a lot more difficult to spot on noisy plots because noisy data causes line segments to cross anyway.

The gradient of a function with a turning point either:

- Starts out positive, decreases to zero, then continues to decrease (maximum value)
- Starts out negative, increases to zero, then continues to increase (minimum value)
- Starts out positive, decreases to zero, then increases again (point of inflexion of an increasing function)
- Starts out negative, increases to zero, then decreases again (point of inflexion of a decreasing function)

Hence the spacings of the y values must do the same.

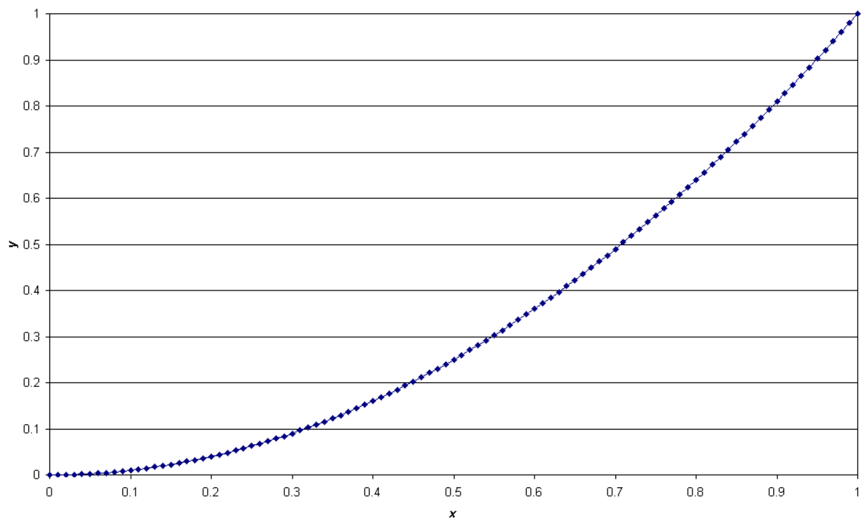

Figure 55: A monotonic function plotted as an X-Y plot.

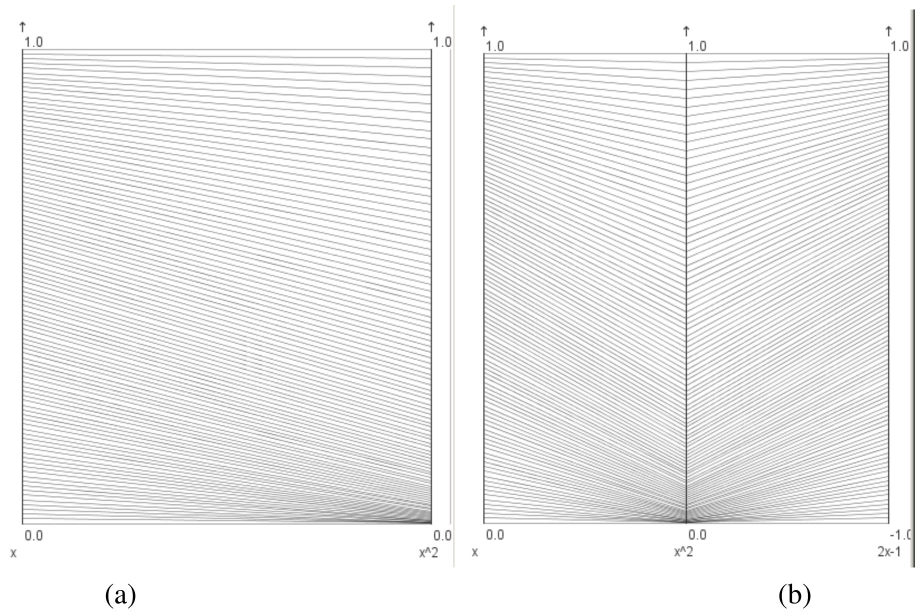

Figure 56: The monotonic function shown in Figure 55, plotted against x as a parallel coordinate plot, (a), and shown between axes of x and $2x - 1$, (b).

First consider the function $x(x - 0.8)$, plotted as an X-Y plot in Figure 57(a). This function has a minimum at $x = 0.4$, as Figure 57(a) shows. The equivalent parallel coordinate plot is shown in Figure 57(b). The spacing of the y values starts off negative (for instance, x_2 is connected to a point lower down the y axis than x_1), increases to zero (hence the cluster of line segments around the bottom of the y axis), and then increases as x does. Hence the plot shows a minimum value.

The plots in Figure 58 show a function with turning points at $x = 0.3$ and $x = 0.8$. Reading from the bottom of the x axis, the y spacing starts off positive, decreases until it passes through zero ($x = 0.3$), reaches a negative minimum, increases and passes through zero again ($x = 0.8$), and is still increasing at $x = 1.0$. This figure is more difficult to interpret than Figure 57(b), because the increased number of crossing line segments make it difficult to trace the behaviour on the right-hand axis.

Inflexion points are turning points where the gradient decreases to zero but does not change sign, so as a parallel coordinate plot the spacing of the y values decreases to zero, but does not change sign so the line segments do not cross. This is illustrated in Figure 59, which shows $(x - 0.5)^3$ as an X-Y plot and a parallel coordinate plot.

7.4 Convergence

Many time-dependent processes feature behaviour that converges to a steady state after a period of time. Examples include hot objects eventually cooling to the temperature of their surroundings and steady state motion of a body after an impulse has been applied. This feature means that it is useful to be able to identify convergence in plots.

Figure 60 shows a function that converges to a limit ($y = 0.1 + 0.9e^{-10x}$) as an X-Y plot and as two parallel coordinate plots, one with the y axis reversed. The gradient of an X-Y plot that converges to a single value tends towards zero, so the line segments in a parallel coordinate plot tend towards zero spacing on the y axis. This is clearer in the axis-reversed plot.

Figure 60 has similarities to the monotonic function shown in Figure 56(b) (in part because the function in Figure 60 is monotonic decreasing). The similarity is also because the function in Figure 56(b) initially has a gradient of zero. Comparison between the two plots shows that Figure 60 has more line segments converging to a single value than the other plot, indicating that its gradient is zero for longer, suggesting a convergence.

In some cases, functions converge to linear behaviour. For example, an object dropped from a great height has a non-linear displacement versus time behaviour initially, but when it reaches its terminal velocity its displacement is a linear function of time. The gradient of an X-Y plot of such a function becomes constant eventually, so it is expected that the vertical spacing between points in a parallel coordinate plot would also become constant. A function that converges to a straight line, $y = x - e^{\frac{(x-0.1)^2}{0.01}}$ is shown in Figure 61. Figure 61(c) has a

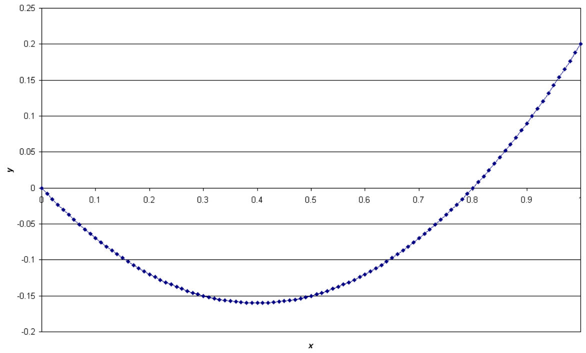

(a)

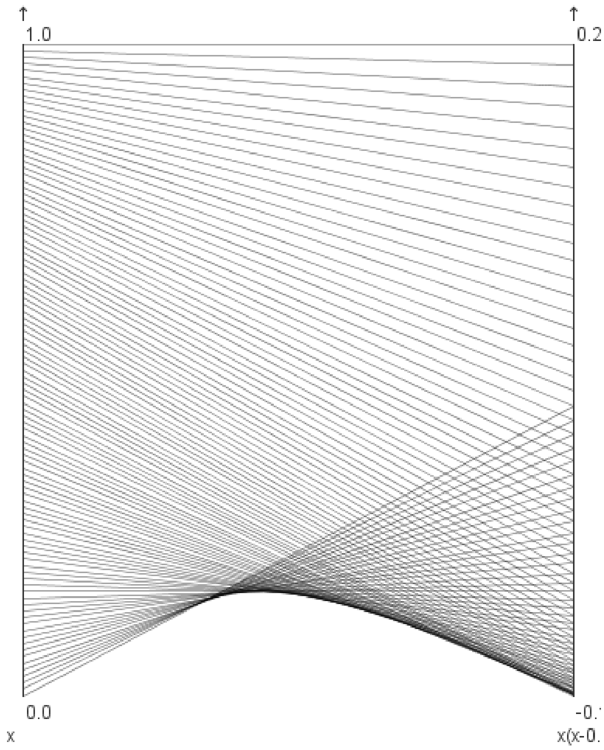

(b)

Figure 57: A function with a single turning point plotted as an X-Y plot, (a), and a parallel coordinate plot, (b).

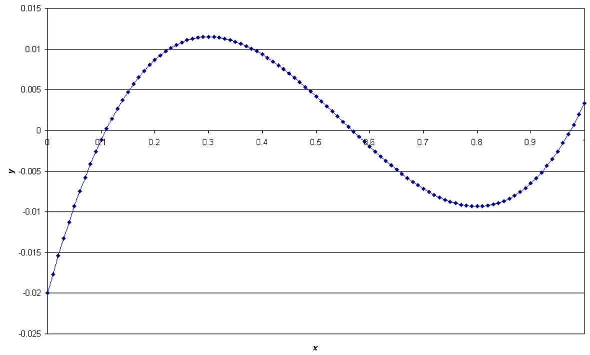

(a)

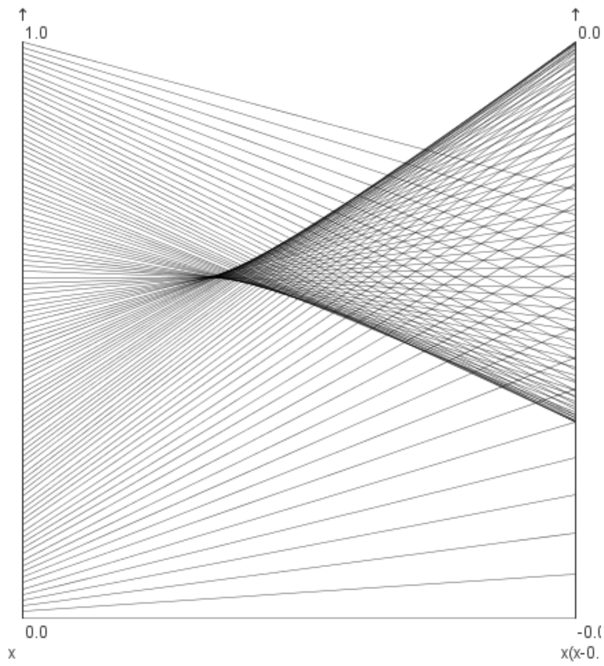

(b)

Figure 58: A function with multiple turning points plotted as an X-Y plot, (a), and a parallel coordinate plot, (b).

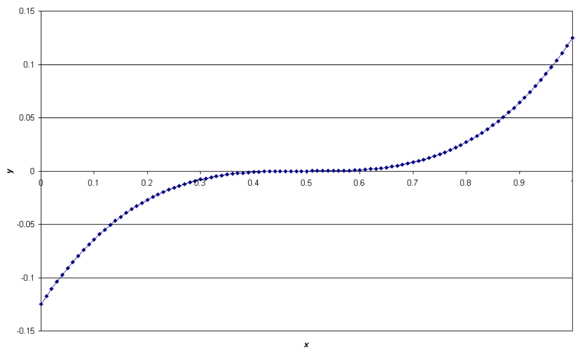

(a)

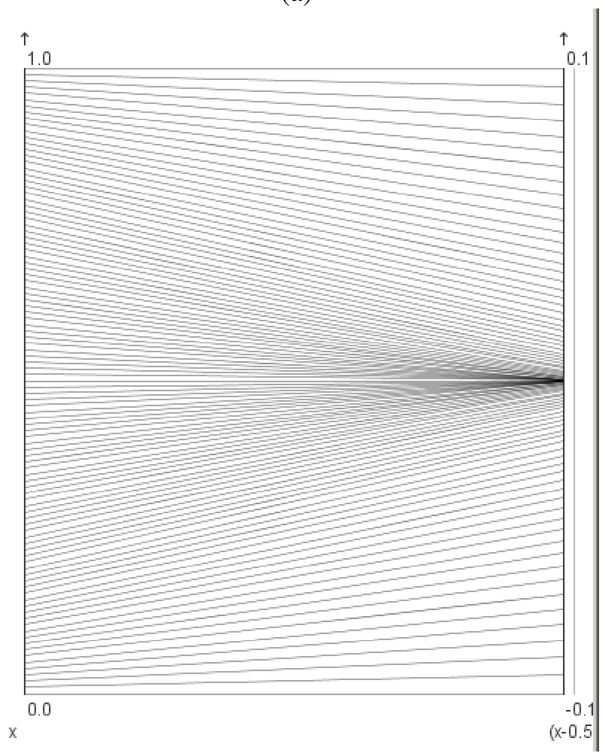

(b)

Figure 59: A function with an inflexion point plotted as an X-Y plot, (a), and a parallel coordinate plot, (b).

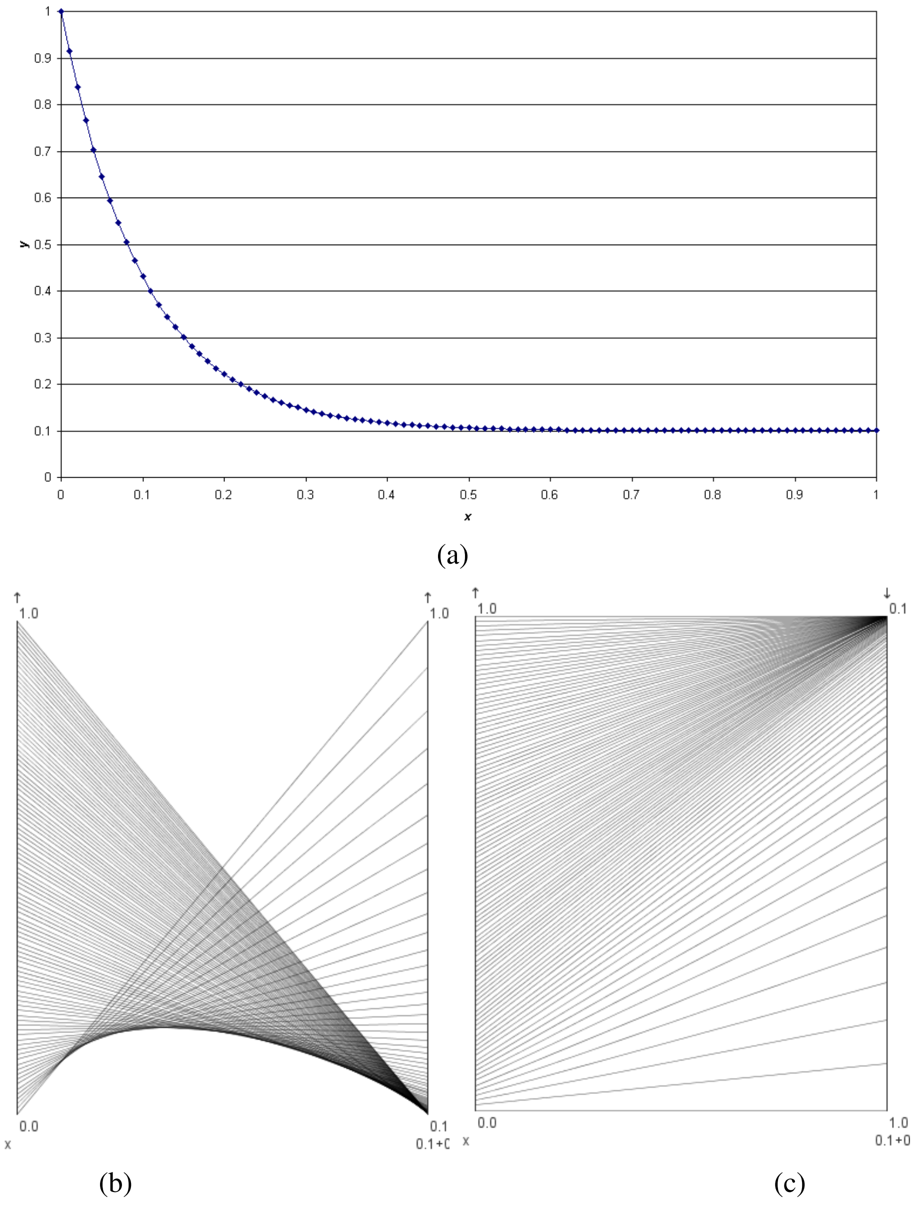

Figure 60: A function that converges to a limit shown as an X-Y plot and two parallel coordinate plots.

rescaled y axis in order to demonstrate the convergence to linear behaviour more effectively: the majority of the line segments in the plot are parallel, as expected.

7.5 Summary

- If the points on one axis in a pairwise comparison are equally spaced, the spacing between the points on the other axis is indicative of the gradient of an X-Y plot of the data being examined.
- Linear relationships are characterised by parallel line segments.
- Monotonic relationships are characterised by non-crossing line segments.
- Turning points can be identified by a spacing between points that reaches zero and then leaves it again.
- Convergence to a single value can be identified by the spacing between points reaching zero and staying there.
- Convergence to a line is indicated if the axes can be rescaled such that the line segments in the convergent region are parallel.

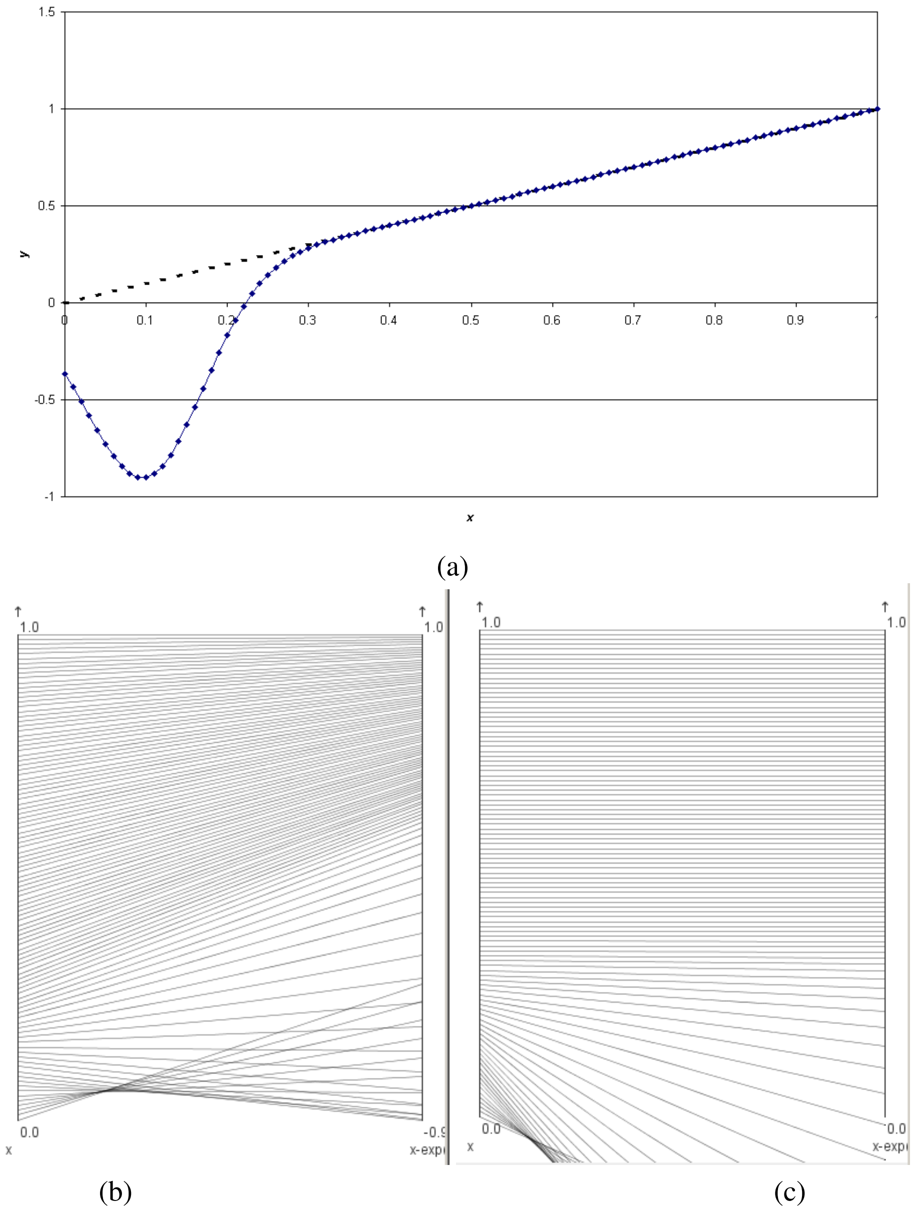

Figure 61: A function that converges to a line shown as an X-Y plot and two parallel coordinate plots.