Graphical Analysis for a Large Designed Experiment

Anne E. Freeny
James M. Landwehr

AT&T Bell Laboratories
Murray Hill, New Jersey 07974

Graphical methods are often used as exploratory tools in analyses of large data sets. However, their use for exploratory purposes in the context of designed experiments is less extensive, perhaps because most experiments produce relatively small amounts of data. In this paper, we show how graphical methods exploiting the structure in a design can be used to reveal interesting patterns in the data and highlight effects not necessarily associated with factors in the design. The methods are presented in the context of an analysis of data from a large experiment (over 3000 observations) designed to investigate the causes of solder defects on circuit boards.

KEY WORDS: Data analysis; Industrial experimentation; Graphical methods; Surface mount soldering.

1. Introduction

Graphical analyses have become accepted as necessary, powerful tools that are an important part of data analysis methodology for many types of statistical problems. For problems where the questions to be answered are not cleanly formulated or where there is much data, various plots that are primarily exploratory in nature are widely used. For problems involving data from designed experiments, where typically there are specific questions and the analysis includes an analysis of variance based on the design, plots are used initially to examine characteristics such as the behavior of the response against time or against each factor. For design of experiment problems, however, additional specialized exploratory plots to be used *prior* to performing the analysis of variance seem not to have received sufficient attention.

This paper shows how specialized, exploratory plots can be used in such problems by giving a case study of the data analysis from a large industrial experiment. The experiment was designed so that relatively straightforward analysis of variance methods, or possibly Poisson regression, should be appropriate for evaluating the results. Graphical analyses show, however, important effects that appear to be unrelated to the design factors. Subsequent analysis uses engineering knowledge and further graphical displays before it makes sense to use statistical models similar to those originally intended, but now using only a subset of the data.

For data from designed experiments, the specialized graphical displays that have been most widely used and discussed are those constructed *after* fitting an analysis of variance model, in particular standard plots

of residuals and full- or half-normal plots of fitted effects (e.g. Daniel, 1976). Additional displays have been proposed that use the fitted model in various ways. Johnson and Tukey (1987) suggest alternative ways of half-normal plotting, while Nair and Pregibon (1986) supplement half-normal plots in the context of robust parameter design experiments with location and dispersion effects. Displays of effects for the various factor levels along with the residuals are used by Hoaglin et. al. (1991) in side-by-side plots, which are similar in spirit to design plots used by Freeny and Landwehr (1990) and by Chambers et. al. (1992); a related display that also includes interaction effects is given by Kenett and Vogel (1991). Multiple comparison results from an analysis of variance can be displayed graphically, and suggestions have been offered by Andrews et. al. (1980) and Sall (1992). In addition, the results of a Bayesian analysis that corresponds roughly to half-normal plotting is displayed graphically by Box and Meyer (1986, 1987).

Here iterative data analysis using graphical methods that are relatively simple, yet specially adapted to the problem and applied *before* the modeling stage of analysis, helps to uncover unexpected effects and convey the results convincingly to both statistician and engineer. The overall message is that important results in design of experiment problems can easily be missed without using graphics in an exploratory way. While we hardly claim that our specific sequence and choice of plots is always appropriate for analyzing designed experiments, we do feel that the approach and ideas illustrated through this case study have broad applicability. Joiner (1977) and King (1992) have also discussed examples where a graphical analysis approach uncovers effects unrelated to the design factors. In a chemometric application, Joiner finds examples of data anomalies and time trends which would contaminate a standard analysis of variance. King describes a situation in which uncontrolled variability in the experimental units was found to be more important than any of the design factors.

The experiment discussed here concerns optimizing a new soldering process for assembling electronic circuit boards. Section 2 describes the problem background and experimental design. After demonstrating in Section 3 that straightforward analysis of variance gives confusing and misleading information, Section 4 presents a series of graphical displays which help to identify patterns of defects that seem to be unrelated to the design factors. We discuss the causes of these defects and why we remove data affected by them. Section 5 then shows how the remaining data can be used to evaluate the effects of the design factors, and

the final section addresses some general issues.

2. Problem Background and Experimental Design

Attaching small electronic devices to a circuit board (sometimes called a printed wiring board) with surface mount soldering requires three operations: application of solder paste on the board, placement of the devices, and reflow of the solder paste. Boards are assembled on a line consisting of three machines, one for each operation. The board first moves into the stencil printer, where it is vary carefully positioned underneath a thin stencil containing many small holes. A squeegee moves across the stencil pushing solder paste through the stencil holes onto pads on the circuit board to which the device leads will be soldered. Experimental factors which can be varied here include the following: properties of the stencil such as material, thickness, and hole size and shape relative to the pads; properties of the squeegee movement such as speed, pressure, and whether one or two passes are made across the stencil; properties of the solder paste; height of the stencil above the board; and frequency of stencil cleaning.

Next the prepared board goes to the device placement machine, which is programmed to move devices of different sizes from feeder tubes to their proper position on the board. Problems which can occur here involve mostly bad parts, which must be screened out before placement; and calibration of the machine, to avoid misplacement which results in the leads of the devices not being correctly centered on the pads.

The board then passes through a long oven which heats the board to a temperature that reflows the solder paste so that it makes the electrical and physical connections of the leads to the pads. Experimental factors include the oven temperature, which can be controlled in each of five sections, and the speed of the board through the oven.

The overall goal of the experimentation is to determine conditions that will give low solder defect rates. The devices soldered in this experiment are substantially smaller than similar ones which were previously used in production. (Smaller devices permit building more electronic functions into a given board.) The factors in the experiment can be thought of as being of two types: physical design factors concerning properties of the circuit board, such as pad dimensions, that can be varied at the circuit board design stage but which are fixed from the perspective of the manufacturing assembly factory; and manufacturing process

factors, such as type of solder paste and machine parameters, which can be varied during assembly. A test circuit board was designed for the experiments as a surrogate for production boards in order to vary physical design, or layout, factors efficiently. Earlier work pointed to the stencil printing machine as the largest source of defects in the process, and separate "set-up" experiments attempted to determine good settings of machine parameters, which were held constant in this experiment. Previous experiments had also narrowed the reasonable range of board layout factors.

The experiment is a classic split-plot situation where some factors vary within boards (split-plot) and others vary among boards (whole plots). The test circuit boards were designed to have 16 areas each containing one device with 100 leads and one device with 132 leads. They arranged in a 4 by 4 grid as shown in Table 1.

TABLE 1. Area Arrangement for Test Circuit Board

FRONT (connector end)			
13	9	5	1
14	10	6	2
15	11	7	3
16	12	8	4

The factors varying over areas within a board are listed in Table 2 and defined as follows. Pad Width and Pad Length give the size of the pads on the circuit board to which the leads are soldered. The abbreviations given for the Geography factor refer to middle Front-Back (areas 5, 8, 9, 12), middle Left-Right (areas 2, 3, 14, 15), corners (areas 1, 4, 13, 16), and center (areas 6, 7, 10, 11). The Geography factor is included in the design because some of the previous experiments suggested that the device position on the board might have some effect on solderability. (In some situations, Geography could be treated as a noise factor in a robust parameter design context. However, here the engineers preferred to consider it in the same way as the other factors.) Device Spacing is the distance between the two adjacent devices in the same area. Stencil Opening refers to the configuration of the stencil opening over each pad. Metal-to-Metal distance refers to the distance from the outer edges of pads on opposite sides of the device. The numerical values of the factor levels given here have been changed from those actually used in the experiment, but the Geography factor and the overall plan of the experiment are preserved.

TABLE 2. Factors Varied Within Each Circuit Board

Factor	Short Name	No of Levels	Levels	Factor Type
Pad Width	PW	4	10, 20, 30, 40	Layout
Geography	Geo	4	mF-B, mL-R, cor, cen	Layout
Device Spacing	DS	2	600, 800	Layout
Pad Length	PL	2	100, 150	Layout
Stencil Opening	SO	2	shape-1, shape-2	Process
Metal-to-Metal	MM	2	1000, 1500	Layout

The experimental design for the 16 areas is given in Table 3. This design was obtained from a 16-run orthogonal array; two groups of three columns were converted to two four-level factors, and four of the remaining columns were chosen as two-level factors so as to give orthogonal estimates of all main effects.

TABLE 3. Factor Levels Assigned to Each Area

Area	PW	Geo	SO	PL	DS	MM
1	10	cor	shape-2	150	800	1000
2	30	mL-R	shape-2	100	800	1000
3	20	mL-R	shape-1	150	800	1000
4	40	cor	shape-1	100	800	1000
5	40	mF-B	shape-2	150	600	1000
6	20	cen	shape-2	100	600	1000
7	30	cen	shape-1	150	600	1000
8	10	mF-B	shape-1	100	600	1000
9	30	mF-B	shape-2	100	800	1500
10	40	cen	shape-1	100	800	1500
11	10	cen	shape-2	150	800	1500
12	20	mF-B	shape-1	150	800	1500
13	20	cor	shape-2	100	600	1500
14	10	mL-R	shape-1	100	600	1500
15	40	mL-R	shape-2	150	600	1500
16	30	cor	shape-1	150	600	1500

The experiment was run in two lots of 48 boards each, one right after the other, and the factors varying among boards are listed in Table 4. The Stencil Thickness values of 5 and 10 mils are not the actual numerical values used. The Board Solder Mask Type represents the type of surface coating applied to the entire board, and the boards were alternated in groups of 12. The stencil was cleaned after every 12 boards, and Board-after-Cleaning gives the corresponding board sequence number (1 to 12). While Stencil Thickness was confounded with lots, we believe that there is no separate lot effect and the division into lots was only for operational convenience.

TABLE 4. Factors Varied Among Circuit Boards

Factor	Short Name	No of Levels	Levels	Factor Type
Stencil Thickness	ST	2	5, 10	Process
Board Solder Mask Type	MT	2	type-1, type-2	Layout
Board-after-Cleaning	BAC	12	1 to 12	Process

Visual observations and electrical test measurements can be made at various points in the assembly process. Here we analyze two responses. The first is number of cross-solders per device (CS), which is determined at the end of the assembly process by electrical in-circuit tests which measure short circuits. The devices were constructed with a "daisy chain" arrangement of circuitry connecting certain adjacent leads so that electronic testing of solder connections could be done, but individual values for all adjacent leads on a device cannot be measured. Instead, the measurement for the device with 100 leads will register from 0 to 15 possible defects; for the device with 132 leads, 0 to 5. The second response is whether or not a part was misplaced, as determined by visual inspection.

3. Analysis of Variance

Earlier experiments showed that the 100- and 132-lead devices soldered similarly. Even though the maximum possible defect values for the two types of devices differ, typically there would not be a large number of defects on a given device. Thus, as a measure of the number of defects in each area of a board, we added the counts of cross-solder defects on both devices. To begin, it appeared that 1267 of the 1536 areas had at least one defect.

Consider an analysis of variance appropriate for this design. With count data the square root transformation is a reasonable initial choice, and the design suggest a model including main effects for all factors and two-factor interactions for the factors that vary among boards. There are two error strata, depending on whether factors vary within boards or among boards.

Table 5 gives this analysis of variance. It appears that the process factors Stencil Thickness and Stencil Opening are "significant," and the layout factor Geography is nearly significant. None of the interactions appear to be important. The estimate of the within-board error variance is less than half of the estimate of the among-board error variance.

TABLE 5. Analysis of Variance sqrt(CS), Both Devices

Error: Among Boards					
	Df	Sum of Sq	Mean Sq	F Value	Pr(F)
ST	1	10.70	10.70	4.95	0.03
MT	1	1.06	1.06	0.49	0.49
BAC	11	27.85	2.53	1.17	0.33
ST:MT	1	1.18	1.18	0.55	0.46
ST:BAC	11	13.38	1.22	0.56	0.85
BAC:MT	11	15.98	1.45	0.67	0.76
Residuals	59	127.59	2.16		
Error: Within Boards					
	Df	Sum of Sq	Mean Sq	F Value	Pr(F)
PW	3	0.44	0.15	0.17	0.91
Geo	3	6.12	2.04	2.44	0.06
SO	1	4.60	4.60	5.51	0.02
PL	1	0.01	0.01	0.02	0.90
DS	1	0.19	0.19	0.22	0.64
MM	1	1.87	1.87	2.24	0.13
Residuals	1430	1193.73	0.83		

Figure 1 shows four graphical summaries of results from this analysis of variance. Figure 1a is a half-normal quantile-quantile (Q-Q) plot of the among-board effects, and Figure 1b is a half-normal Q-Q plot of the within-board effects. Both plots show good straight-line configurations. No effects appear to stand out, although we would have expected some from the analysis of variance table. This surprising result raises some suspicions about the validity (not the numerical accuracy) of the analysis of variance. Looking next at the bottom of Figure 1, panel 1c is a normal plot of the residuals and panel 1d plots the residuals against the square root of the observed values. These plots raise further suspicions. Figure 1c shows a decidedly non-normal configuration in the lower end. The separate vertical segments in Figure 1d result from the observed response (recall that this analysis uses square root of defects) being discrete, and the very wide range of residuals for each observed value emphasizes that the fit is very poor. Adding all estimable interaction terms and fitting a Poisson model directly (not shown) does not give a tighter fit nor better behaved residual plots. Clearly we need more detailed analyses to find out what is going on in these data.

4. Exploratory Displays for Data Investigation

Since we have reason to question the usefulness of the overall anova model and analysis, we choose to step back and look at the raw data in more detail through several different plots. Figure 2a shows box plots

of CS defects for devices with 100 leads for all 48 boards in lot 1, and Figure 2b gives the corresponding plot for devices with 132 leads. Each box is calculated from 16 data values and shows the distribution of cross-solder defects over the 16 areas in a board. It is immediately obvious that the cross-solder behavior is not the same for the two types of devices. For devices with 100 leads, the configuration is a full box for every board, so the median number of defects per area is greater than zero for each board. For devices with 132 leads, many boxes have degenerated to a line at 0, indicating that there are no defects in any area. Other boards are represented by a line at 0 with one or more outliers plotted (with potential overstrikes) at 1, 2 or 3, so at least 3/4 of the defect counts are 0 for these boards. A few boards show a half box, indicating that at least half of the 132-lead devices have no defects. The corresponding plot from lot 2 (not shown) exhibits similar results for the two types of devices. Thus, there are clearly many more defects for devices with 100 leads consistently across the boards in the experiment.

Such different soldering behavior for the two types of devices was not expected, so we needed to learn why it occurred. Investigation revealed that the two types of devices were made in different lots with different processes. Testing some of the devices not used in the board assembly revealed that for the devices with 100 leads there were internal shorts across the leads in many cases, with an average of about 4 shorts per device. However, because of the daisy chain testing of groups of leads together, the defect count would typically be less than 4. The devices with 132 leads were functional. Therefore, electrical test measurements for the 100-lead devices cannot be associated with the ability of the soldering process or the effects of the design factors. We must discard the CS data from the devices with 100 leads for the purpose of learning about the design factors. (It is also quite valuable, of course, to learn that the process used to manufacture the 100-lead devices was faulty.)

For the devices with 132 leads, the total number of CS defects is 162 over all 96 boards in both lots. Figure 2b does not give detailed information concerning how these defects vary from board to board in lot 1, although there is an indication of more than a typical number of defects on the first three or four boards. We next want to look at the data in a way that can show more detail than do box plots concerning the defect counts from board to board. We also want to examine the defect counts at different levels of aggregation than the whole board.

Since the defect count per device ranges from 0 to only 5, we can achieve a more informative display by using shaded bar plots which show, for each board, the number of devices with each of the possible number of defects. Figures 3a and 3b are shaded bar plots for lots 1 and 2, respectively, with boards plotted in run order. The area in each bar with the different possible gray-scale shadings corresponds to the proportion of devices on each board with the defect counts, as indicated in the key. Bars that are completely white indicate no defects for those boards, and lighter shaded bars have fewer defects than darker bars.

Figure 3a shows that there is a run of four boards with defects at the beginning of lot 1, but no similar run anywhere else. What is the explanation? Referring back to notes taken while the experiment was being run, we find that the process engineer noticed that the placement machine appeared to be misplacing devices in the first few boards in lot 1. The assembly was stopped after the fourth board and the machine was recalibrated. A plot of the number of misplaced parts per board (not shown) corroborates that there were many misplaced parts in the first four boards of lot 1 but very few thereafter. Since misplaced parts generally would cause cross-solders that are not a result of the factors of the soldering process or the physical design, we removed the data from these four boards from consideration as well.

Figures 3a and 3b also clearly show that when there are defects, over both lots except for the first four boards in lot 1, the defects tend to be on every other board. Why is this? At this point, we don't know and would eventually like to find possible reasons.

In order to understand how the within board factors (Table 2) affect the results, we need displays relating CS defects to these factors. Shaded bar plots of the defects for each of the 16 areas across all boards except for the first four in lot 1, and for all boards in lot 2 (these plots are not shown), indicate that the pattern of defects was the same in both lots. Thus, we combine the defect data over all boards and display shaded bar plots of the 16 areas in Figure 4, which is arranged corresponding to the layout of areas on a board as given in Table 1. It is clear that areas 1, 5, 9, and 13 -- which are all in Row 1 -- have more defects than the other areas. Row 1 does not, however, directly correspond to any of the within-board design factors. Table 3 shows that the factors PW, DS, PL, and MM are all balanced across these four areas; these areas have the same level of SO (shape-2) and include two of the four levels of Geo.

One might suspect that the common SO level of shape-2 is the cause for the relatively poor performance of areas 1, 5, 9 and 13. Areas 2, 6, 11 and 15 also have SO level of shape-2 and the same balance across the other design factors; therefore, it is worthwhile to compare the performance of these areas with areas 1, 5, 9, and 13 using Figure 4. Areas 2, 6, 11 and 15 do have a few more defects than the other eight areas not in the first row, but they do not have nearly as many defects as the devices in the first row. Thus, a first analysis suggests that it is unlikely that poor performance at SO level shape-2 is the entire explanation. Similarly, the bottom row has the same levels of the Geo factor as the first row but does not have nearly as poor performance.

We further investigate the defect data displayed in Figure 4 by plotting defects against the levels of both the within-board and among-board factors. Figure 5 shows shaded bar plots for each of the factors. The heights of the bars for the among-board factors are not quite the same since data from the first four boards from lot 1 are omitted. The difference in heights helps show the lack of balance in the among-board part of the design thus introduced. Although there are relatively few defects so that most of the plotted bars are white, Figure 5 still suggests some indications of "factor effects." The largest effects seem to be for SO and Geo for the within-board factors, and MT and possibly BAC for the factors that vary between boards.

The large amount of white space in Figure 5 makes it difficult to make precise visual comparisons of the effects of different levels of the factors. Displaying this data in another way helps in this regard. Figure 6 shows the proportion of the devices for each level of each factor that have a defect count greater than zero. Thus, we combine defect counts of 1, 2, 3, 4, and 5, and display their total proportion in a standard bar chart or histogram. Since the maximum proportion shown is 0.15, each plot uses more of the available space to display the actual differences between levels than was the case in Figure 5. A disadvantage is that Figure 6 does not show the places where unbalance occurs.

Figure 6 shows that there is hardly any difference in the proportion of devices with non-zero CS defects across the levels of DS, MM, and ST. The levels of factors PL and PW show slightly more differences, while the strongest differences involve factors SO, Geo, BAC, and MT. The most interesting plot is that for BAC, which clearly shows that there are many more defects for odd-numbered boards than even-numbered ones, except for possibly the first board. This result reinforces an observation from Figure 3.

At this point we have discovered an apparently large effect that seems related to the first row, which does not directly correspond to a design factor, so we do not have any clean interpretation for it. Moreover, there may be effects from some design factors, and it is not clear to what extent they may be separate from the effect associated with the first row. We would like to understand the data better, and a reasonable next step is to examine interactions between some of the factors that seem most important. As part of this investigation we define a 'pseudo-factor' called Row, with level 1 representing the top row of areas 1, 5, 9, and 13. This factor is balanced with respect to the among-board factors, though as we have seen it is not balanced with respect to all the within-board factors and thus requires careful interpretation.

Because of our specific experimental design, interactions can be constructed only for pairs of among-board factors and among-board by within-board combinations, but not for pairs of within-board factors. The factors that seem to be most important are SO, Geo, Row, BAC and MT. Thus, we construct interaction plots for all permissible pairs of these factors and show them in Figure 7. The proportion of the number of devices with cross-solders is plotted on the y-axis for all plots. These plots are more complex than the previous ones, and examining them requires some time and care.

The four plots with Board after Cleaning on the x-axis show that there are strong interactions of BAC with each of these within-board factors, and the odd-numbered boards perform worse. The plots of SO against BAC and Row against BAC appear to have the strongest interactions. There are many defects for row 1 and odd-numbered boards (except the first), and many fewer defects for the other rows; similarly, the level shape-2 of SO has many defects for the odd-numbered boards. (Recall that row 1 contains only SO level shape-2.)

Thus, there appears to be an effect of the leading edge (Row 1) on odd-numbered boards. The only process parameter which differs on every other board is the direction the squeegee travels when applying solder paste in the stencil printer. For the odd-numbered boards the squeegee moves from left to right, and devices in the four areas 1, 5, 9, 13 are the first to have solder paste deposited. For even-numbered boards, they are the last.

Among the plots with MT on the x-axis, the plot showing Row shows the strongest indication of interaction. Here boards with level type-2 of MT perform better in row 1 than do those with level type-1.

Since pairwise plots of BAC, MT, and Row all indicate two-factor interactions, we should examine their three-factor interaction, which is shown in Figure 8. The Row 1 odd-even effect is much stronger for boards of type-1 MT, although there is some indication of the odd-even effect for type-2 MT as well.

For purposes of evaluating the design factors for this experiment, we thus decide to remove the data from areas 1, 5, 9, and 13 for odd-numbered boards. There is a strong indication that the many defects in these places were caused by some effect different from the design factors. In addition, of course, we need to investigate possible engineering reasons for these defects and ways to resolve whatever the problem is.

We went back to the process and design engineers, who at first had no idea why this problem should occur. Eventually, however, they extended the squeegee travel distance at the end of the right-to-left wipes so that the solder paste would be pushed further from the edges of the first row areas before the odd-numbered wipes. Within the areas in the first row, the 132-lead devices are at the edge of the board while the 100-lead devices are in the parts of these areas away from the edge, so on the left-to-right (odd-numbered) squeegee passes the 132-lead devices are the first to receive the solder paste. The idea is that the solder paste may need to be "rolling smoothly" along the board in order to be deposited adequately through the stencil openings, and it might not do this at the very beginning of a wipe. Moreover, on the first pass after cleaning the stencil (BAC of 1), the squeegee starts farther to the left, and the defect counts for these boards were not as high as for the later odd-numbered boards. Thus, this change and explanations are plausible and do agree with the results found in the data analysis. The engineers also changed to a wider squeegee. Work subsequent to this experiment showed that these changes resolved the problem.

5. Assessing Design Factor Effects

The results from our sequence of exploratory plots, suggest that special problems affected the data from the areas 1, 5, 9, and 13 on odd-numbered boards, as well as the first four boards in lot 1, and all the devices with 100 leads. Thus these data should be removed before evaluating the effects of the design factors on cross-solders. Now, however, only 21 CS defects remain in the data and they are spread over 17 different devices. Originally we considered defect counts on 3,072 different devices, and after the data deletions 1288 different devices remain. In one sense the small proportion of devices with CS defects is a very good

result, as it implies that after removal from the analysis of defects not associated with "routine" operation of the soldering process, the soldering process appears to be working quite well. On the other hand, a major purpose of the experiment was to learn about possible differences in effects from the different levels of the design factors, and with only a few overall defects any possible differences are harder to pin down.

A way to graphically display and assess the remaining data is shown in Figure 9. The number of devices for which we have a defect count is shown as the height of the thin white bars and can be read using the scale at the right of each plot; the proportion of these observations with non-zero defects is shown by the heights of the black bars, which correspond to the scale at the left and which is the same across all pieces of the figure. The differences in heights of the white bars across the levels of a factor show the extent to which the data deletion has created unbalance for each factor. The differences in heights of the black bars show the extent to which there is a difference in defect behavior across the levels of each factor. The lack of balance implies that we must be careful about interpreting these effects. The plot shows that the largest remaining effect is for SO; shape-1 has a proportion of devices with non-zero defects of about 0.005, while Shape-2 has a proportion of 0.024. Of the remaining factors, Geo, PL, and MT show the largest differences in defect rates. For Geo, the first and fourth rows (mF-B and cor) perform a little better than the middle rows (mL-R and cen). A PL level of 150 appears somewhat better than a level of 100, and MT of type-2 looks preferable to type-1. The factor BAC also shows some variability among the defect rates over its levels, but note that there are fewer observations at each level than for the other factors and there is no clear pattern across the levels of BAC.

It is also useful to summarize the 21 CS defects in a table showing the number of defects at each level of each design factor, as in Table 6. This table also shows the attained significance level *P* for each factor from testing the null hypothesis that both levels of the factor give equal defect rates; this test assumes a Poisson distribution of defects and is conditional on observing 21 total defects. These tests are not entirely uncorrelated because of the unbalance resulting from the data deletions.

The factor SO is the only one for which the difference between levels is large enough to be statistically significant. The other factors with differences closest to being statistically significant are Geo, MT, and PL.

TABLE 6. Remaining Defect Counts and Significance Level for Design Factors

Factor Names	Levels	Defect Counts	Sig. Level *P*	Compared
Pad Width	10, 20, 30, 40	3, 5, 7, 6	.38	10+20 vs 30+40
Geography	mF-B, mL-R, cor, cen	3, 10, 2, 6	.12	cor+mF-B vs cen+mL-R
Device Spacing	600, 800	9, 12	.66	
Pad Length	100, 150	14, 7	.19	
Stencil Opening	shape-1, shape-2	4, 17	.0008	
Metal-to-metal	1000, 1500	13, 8	.38	
Mask Type	type-1, type-2	14, 7	.13	
Stencil	5, 10	8, 13	.50	
Board-after-Cleaning	1-3, 4-6, 7-9, 10-12	4, 3, 11, 3	.26	1 to 6 vs 7 to 12

These calculations confirm and make more precise the impressions given by Figure 9. The experimental results suggest that the choice of levels for the other factors can be made according to other considerations, such as cost.

It is of interest to estimate what the production CS defect rate would be under the optimal conditions identified in this experiment. To do this, we assume that the placement machine is calibrated correctly and that the squeegee travel problem on the leading edge of the odd-numbered boards is resolved. In other words, the problems identified that were not related to the factors in the experiment do not reoccur, so we can estimate defect rates using the final data after the indicated parts are deleted. We also assume that an electrically measured CS defect, which because of the daisy-chaining for testing purposes means that there is one or more shorts over several adjacent leads, actually corresponds to only one physical cross-solder for some pair of leads. Pooling over all the remaining data with shape-1, we have an estimated defect rate of 41 ppm, meaning 41 cross-solders per 1,000,000 leads soldered under these conditions.

6. Discussion

We used relatively simple graphical data analysis methods to discover that the major causes of defects in this experiment were not related to the design factors. A combination of careful data analysis along with engineering information enabled us to conclude that most defects were due to non-functional devices, a placement machine miscalibration, and a complicated, completely unexpected problem with the stencil printing machine. Analyzing these data using standard anova techniques and without discovering these problems gives misleading results. However, once these problems are identified and the corresponding data removed, we are able to get information about the effects of the design factors. It turns out that one factor

has a statistically significant effect, several have possible but more marginal effects, and several don't seem to matter much.

Our discussion and development of these displays have been in the context of a specific study, but we feel that the types of displays that proved useful here have broad applicability in design of experiment problems. While displays often have to be tailored to the specific problem, some characteristics useful here and more generally are the following: the graphs focus on raw data values structured in an appropriate way; they show variation as well as means; they use ordering over time; and they can be easily interpreted by both engineers and statisticians.

Our approach also illustrates and reinforces several general principles of statistics that apply to a broad range of experiments. None of the following points is original here, but they are nonetheless important and need to be a conscious part of statistical practice and thinking.

There is no substitute for careful and detailed data analysis, even for an experiment that is well designed and carefully run. Some simple, exploratory plots should always be used early in an analysis, even when a model seems obvious based on the design. It helps to use a sequence of plots, first to explore general characteristics and then to focus on relevant details. Simple plots can be very useful, though they may have to be adapted to the special characteristics and objectives of the specific problem at hand. In summary, data analysis insight may be more important than using complicated modeling techniques to reach valid conclusions.

However, apart from the data analysis methodology, it is equally important to know the physical background for the problem and to take good notes during the experimentation. Simply finding anomalous data and removing them is not adequate, unless there is also some plausible physical reason for doing so that can be checked through subsequent work.

Finally, formal statistical inference is often a very small part of the overall study.

7. Acknowledgements

The task force to introduce this technology was comprised of many people who made important contributions. We single out especially our engineering collaborators Greg Munie and Bob Prazak, without

whose insight and leadership the project would not have succeeded. We also thank Marilyn Becker for assistance in management and analysis of the data.

REFERENCES

Andrews, H. P., Snee, R. D., and Sarner, M. H. (1980), "Graphical Display of Means," *The American Statistician*, 34, 195-199.

Box, G. E. P., and Meyer, R. D. (1986), "An Analysis for Unreplicated Fractional Factorials," *Technometrics*, 28, 11-18.

Box, G. E. P., and Meyer, R. D. (1987), "Analysis of Unreplicated Factorials Allowing for Possibly Faulty Observations," pp 1-12 in *Design, Data, and Analysis*, Ed. C. L. Mallows, John Wiley and Sons, New York.

Chambers, J. M., Freeny, A. E., and Heiberger, R. M. (1992), "Analysis of Variance; Designed Experiments," pp 145-193 in *Statistical Models in S*, Ed. J. M. Chambers and T. Hastie, Wadsworth & Brooks/Cole Pacific Grove, CA.

Daniel, C. (1976), *Applications of Statistics to Industrial Experimentation*, John Wiley and Sons, New York.

Freeny, A. E., and Landwehr, J. M. (1992), "Displays for Data from Large Designed Experiments," pp 117-126 in *Computing Science and Statistics: Proceedings of the 22nd Symposium on the Interface*, Eds. C. Page and R. LePage, Springer-Verlag, New York.

Hoaglin, D. C., Mosteller, F., and Tukey, J. T., Eds. (1991), *Fundamentals of Exploratory Analysis of Variance*, John Wiley and Sons, New York.

Johnson, E. G., and Tukey, J. W. (1987), "Graphical Exploratory Analysis of Variance Illustrated on a Splitting of the Johnson and Tsao Data", pp 171-244 in *Design, Data, and Analysis*, Ed. C. L. Mallows, John Wiley and Sons, New York.

Joiner, B. L. (1977), "Evaluation of Cryogenic Flow Meters: An Example in Non-Standard Experimental Design and Analysis," *Technometrics*, 19, 353-379.

Kenett, R. S., and Vogel, B. M. (1991), "Going Beyond Main-Effect Plots," *Quality Progress,* February 1991, 71-73.

King, J. R. (1992), "Presenting Experimental Data Effectively," *Quality Engineering,* 4, 399-412.

Nair, V. N., and Pregibon, D. (1986), "A Data Analysis Strategy for Quality Engineering Experiments," *AT&T Technical Journal,* 65, 73-84.

Sall, J. (1992), "Graphical Comparison of Means," *Statistical Computing and Statistical Graphics Newsletter,* April 1992, 27-32.

FIGURE CAPTIONS

Figure 1. Summary plots for analysis of variance of transformed counts of cross-solders over both 100- and 132-lead devices: (a) Half-normal Q-Q plots for among-board factors; (b) Half-normal Q-Q plots for within-board factors; (c) Normal Q-Q plot of residuals; (d) Residuals plotted against transformed counts of cross-solders. The model includes main effects and two-factor interactions for factors which vary among boards, and main effects for factors which vary within boards.

Figure 2. Box plots of number of cross-solders on each board in lot 1: (a) 100-lead devices; (b) 132-lead devices. When the box degenerates to a line at 0, there are no defects on that board. Counts plotted individually by circles are beyond 1.5 interquartile ranges of the nearest quartile.

Figure 3. Shaded bar plots of number of cross-solders for 132-lead devices on each board: (a) Lot 1; (b) Lot 2. The area in each bar with different grey-scale shadings corresponds to the proportion of devices with defect counts from 0 to 5 as given by the key. The height of the white space indicates the number of areas with no defects.

Figure 4. Shaded bar plots of number of cross-solders in each area. Boards 1 to 4 in lot 1 are omitted. The arrangement of bars corresponds to the position of the areas on a board as given in Table 1.

Figure 5. Shaded bar plots of number of cross-solders by factors. Boards 1 to 4 in lot 1 are omitted. The different heights of the bars within certain factors is caused by the lack of balance resulting from omitting these boards.

Figure 6. Proportion non-zero cross-solder defect plots by factors, omitting boards 1 to 4 in lot 1.

Figure 7. Interaction plots of proportion non-zero cross-solder defects, omitting boards 1 to 4 in lot 1. The similar appearance of the first four panels is largely due to the confounding of the pseudo-factor Row with the factors SO and Geo.

Figure 8. Three-factor interaction for factors BAC, Row, and MT, displayed by proportion non-zero cross-solder defects vs BAC: (a) For MT = type-1; (b) For MT = type-2. Boards 1 to 4 in lot 1 are omitted.

Figure 9. Enhanced proportion non-zero cross-solder defect plots by factors. Boards 1 to 4 in lot 1 and areas 1, 5, 9, and 13 on odd boards are omitted. Within each factor the proportion of non-zero defects is given by the height of the black bars (left scale) and the number of devices is shown by the height of the white bars (right scale) for each level of the factor. The difference between the heights of the white bars in each subfigure reflects the lack of balance created by data deletion for that factor.

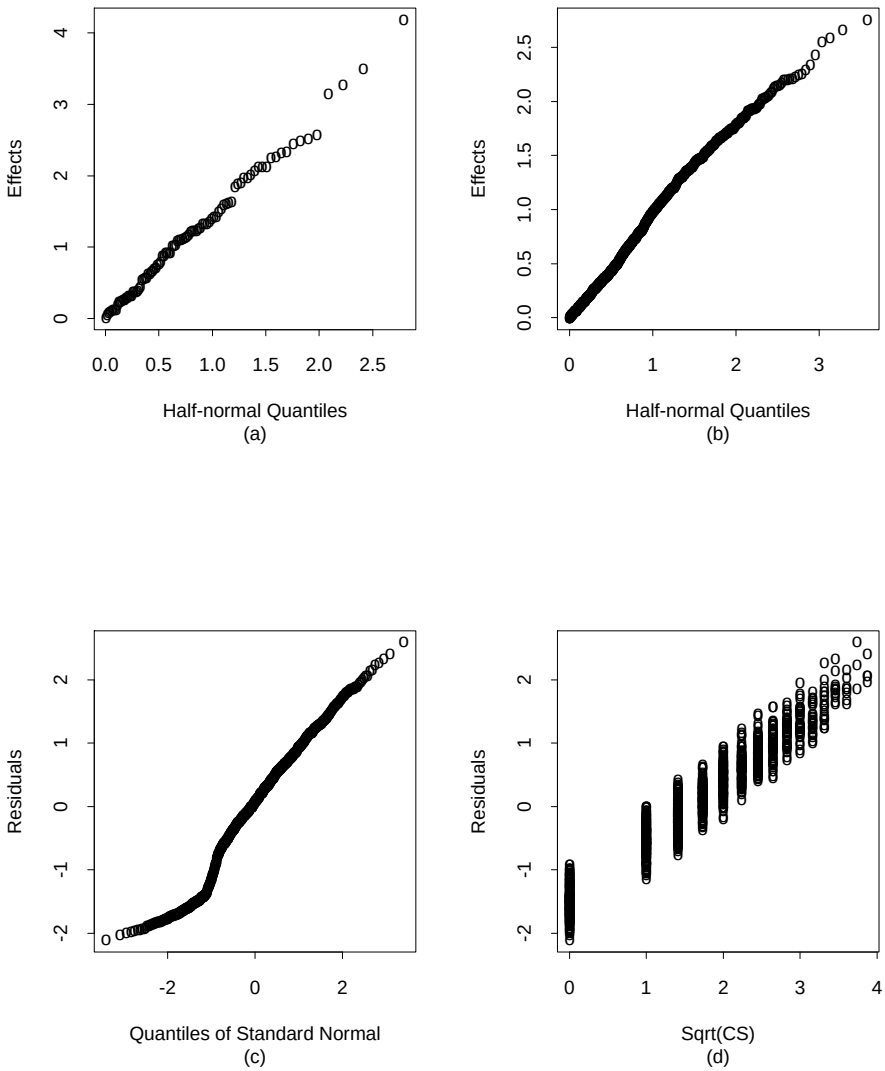
Effects
Half-normal Quantiles
(a)
Effects
Half-normal Quantiles
(b)
Residuals
Quantiles of Standard Normal
(c)
Residuals
Sqrt(CS)
(d)

FIGURE 1

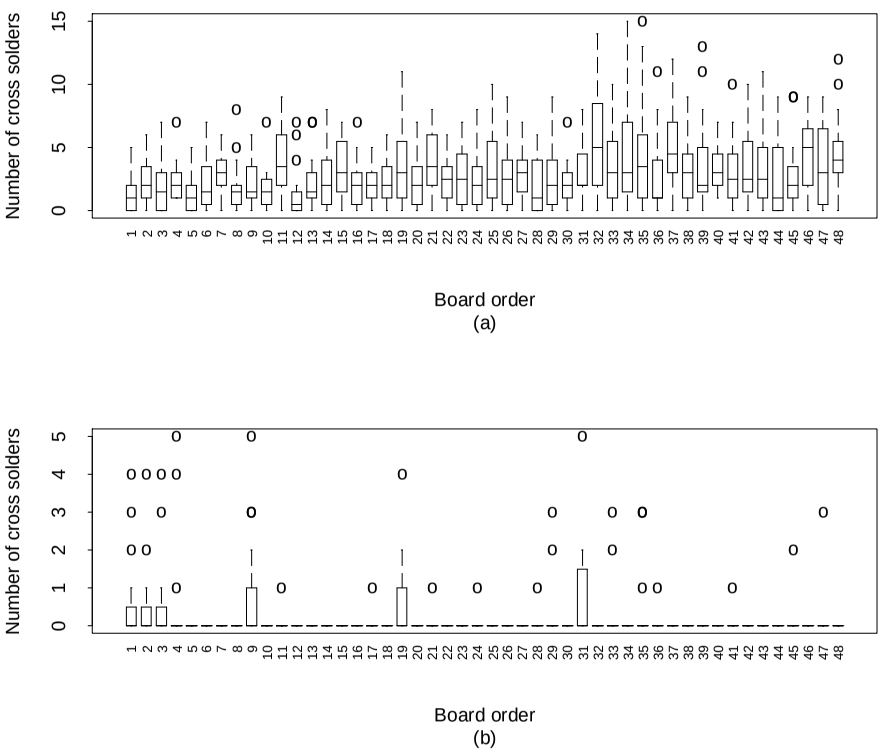
Number of cross solders
0
5
10
15
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48
Board order
(a)
Number of cross solders
0 1 2 3 4 5
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48
Board order
(b)

FIGURE 2

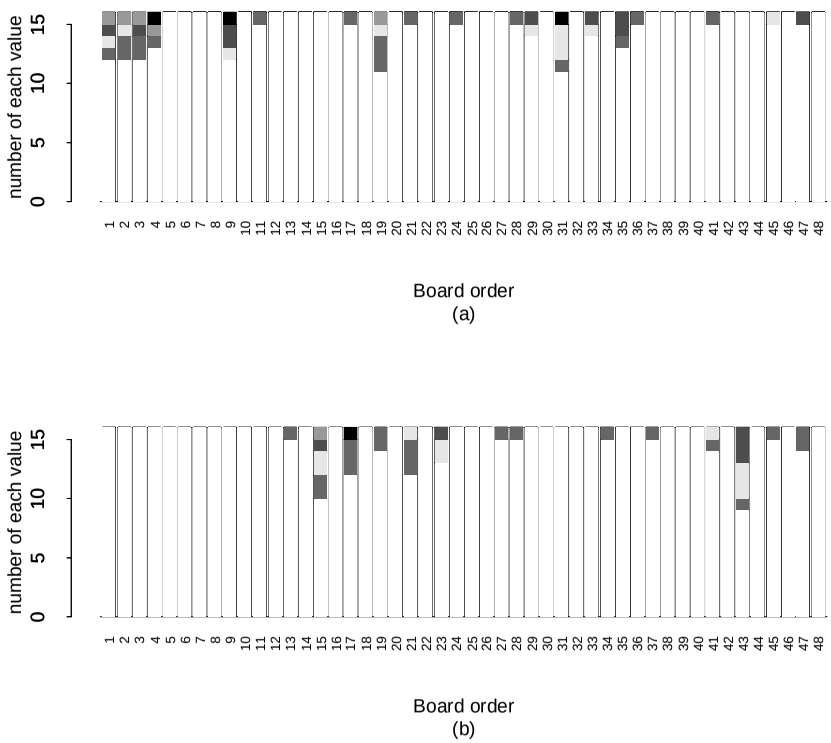
number of each value
0
5
10
15
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48
Board order
(a)
number of each value
0
5
10
15
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48
Board order
(b)

FIGURE 3

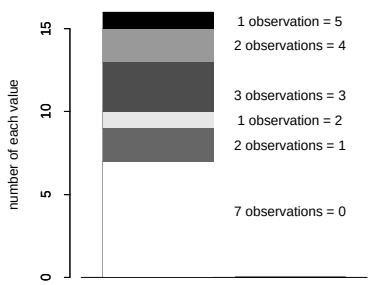

Key for shaded bar - 16 observations, 6 values

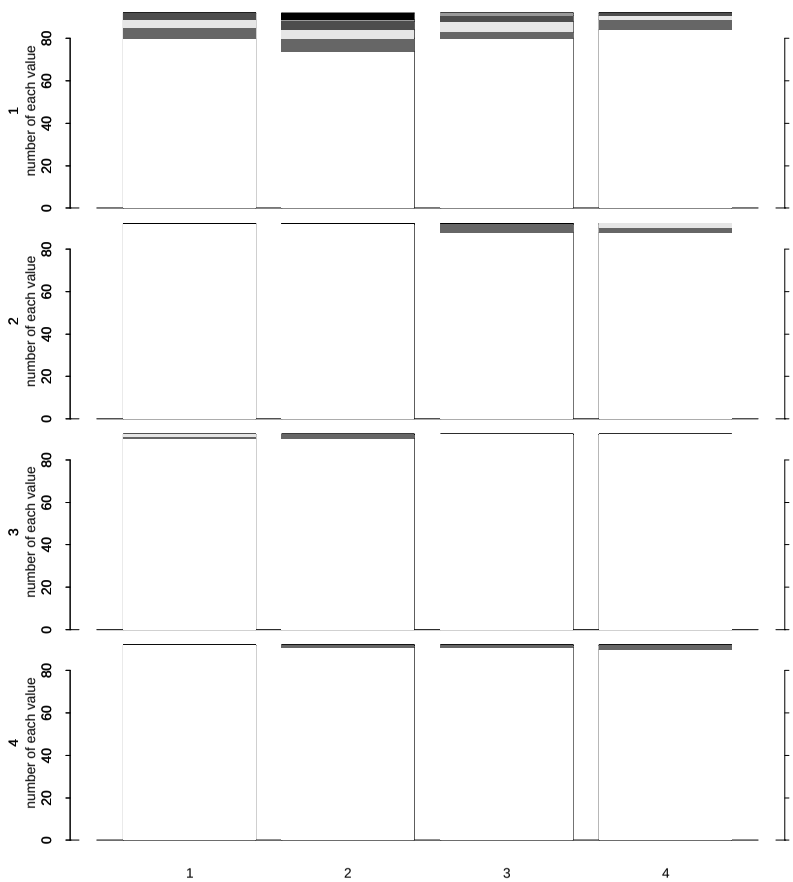
1
number of each value
0
20
40
60
80
2
number of each value
0
20
40
60
80
3
number of each value
0
20
40
60
80
4
number of each value
0
20
40
60
80
1
2
3
4

FIGURE 4

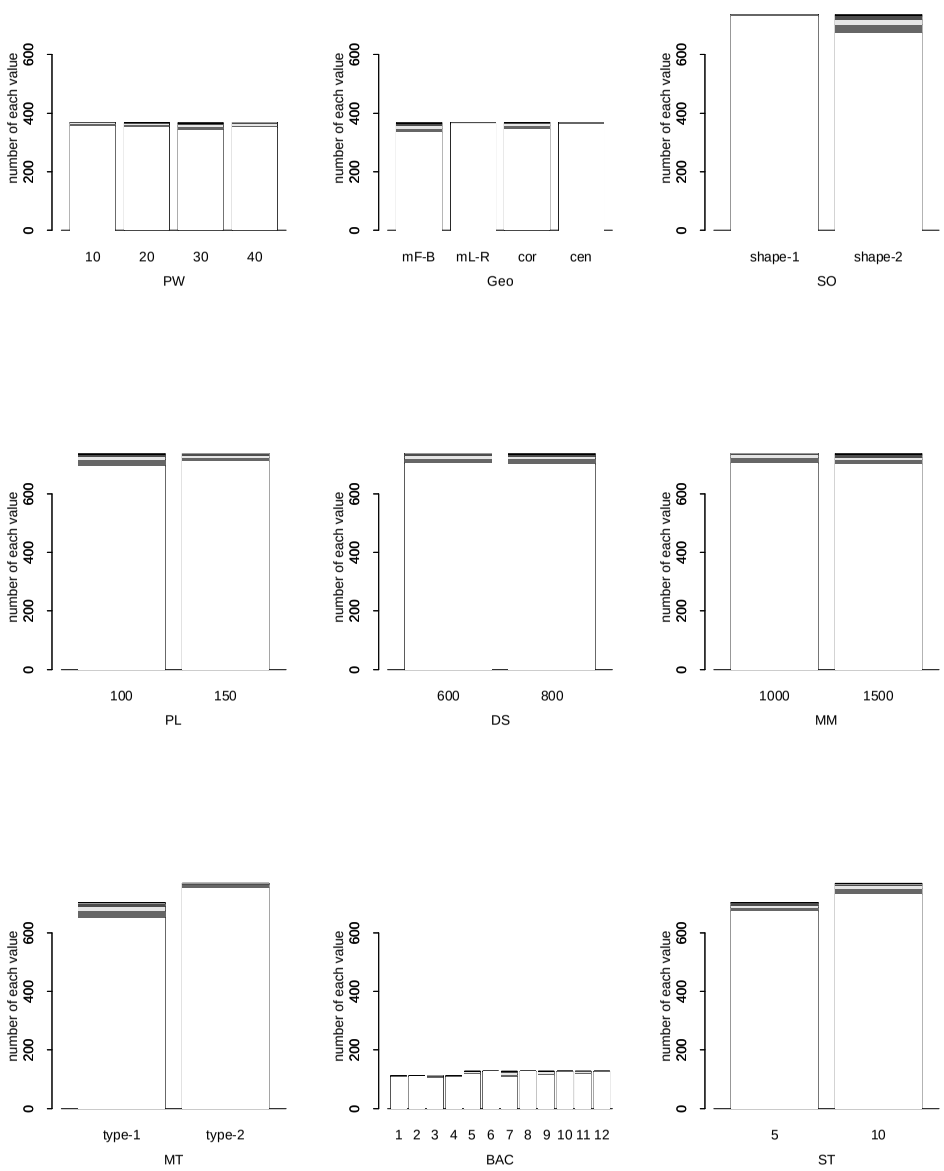
number of each value
0
200
400
600
10
20
30
40
PW
mF-B
mL-R
cor
cen
Geo
shape-1
shape-2
SO
100
150
PL
600
800
DS
1000
1500
MM
type-1
type-2
MT
1 2 3 4 5 6 7 8 9 10 11 12
BAC
5
10
ST

FIGURE 5

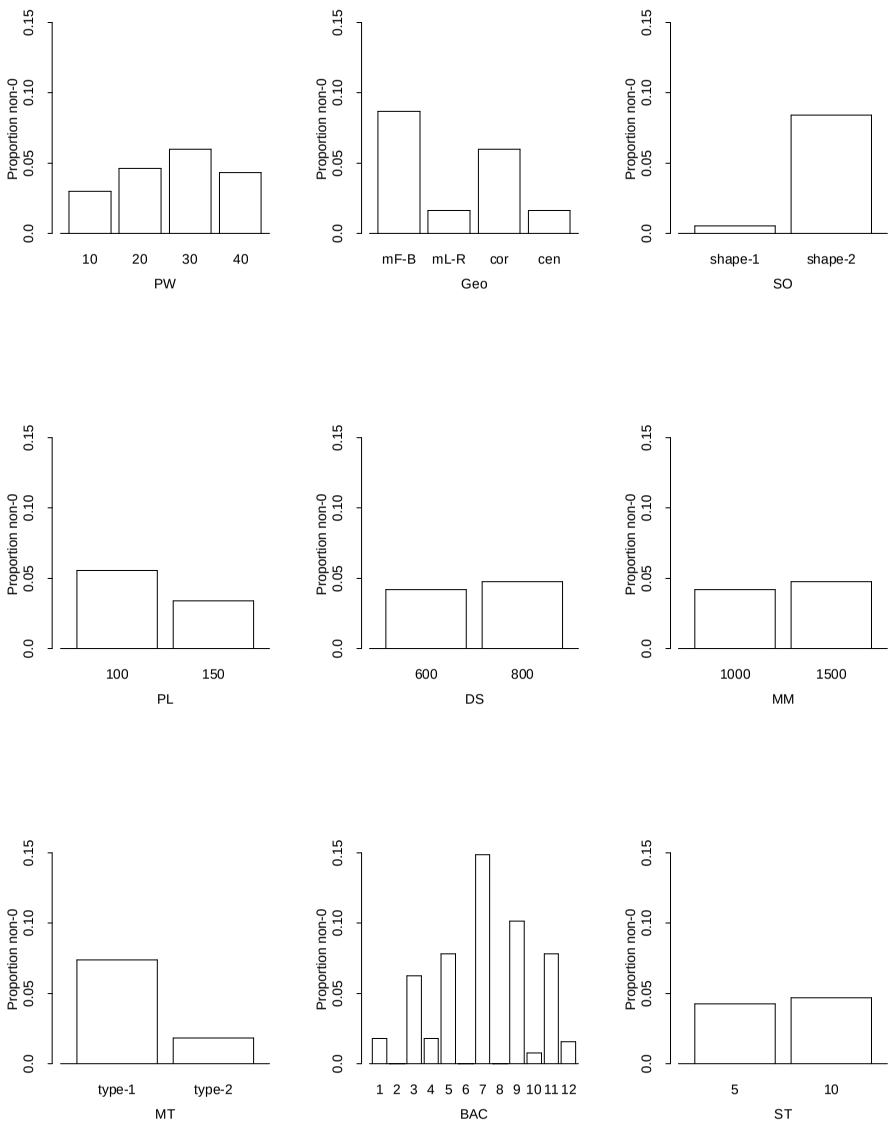
Proportion non-0
0.0
0.05
0.10
0.15
10
20
30
40
PW
mF-B
mL-R
cor
cen
Geo
shape-1
shape-2
SO
100
150
PL
600
800
DS
1000
1500
MM
type-1
type-2
MT
1 2 3 4 5 6 7 8 9 10 11 12
BAC
5
10
ST

FIGURE 6

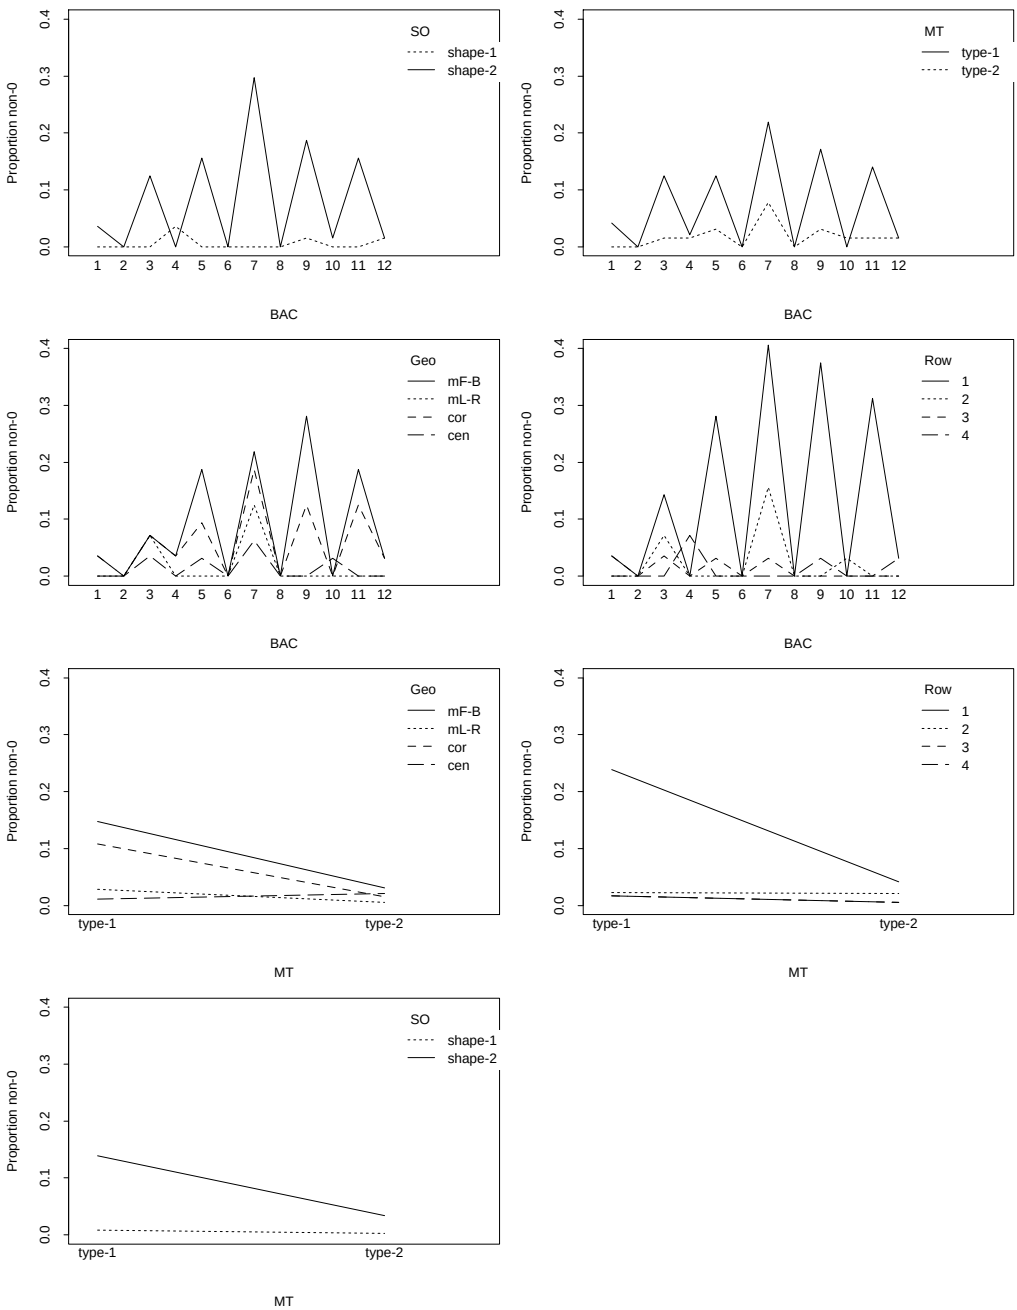
Proportion non-0
0.0
0.1
0.2
0.3
0.4
1 2 3 4 5 6 7 8 9 10 11 12
BAC
SO
shape-1
shape-2
MT
type-1
type-2
Geo
mF-B
mL-R
cor
cen
Row
1
2
3
4
type-1
type-2
MT

FIGURE 7

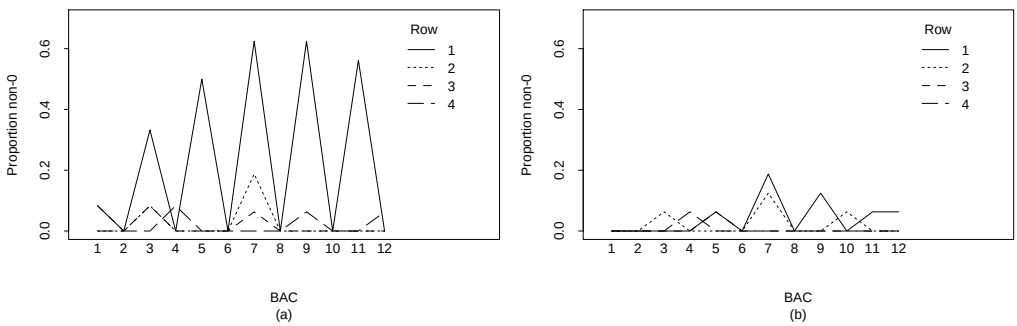
Proportion non-0
0.0
0.2
0.4
0.6
1 2 3 4 5 6 7 8 9 10 11 12
BAC
(a)
Row
1
2
3
4
Proportion non-0
0.0
0.2
0.4
0.6
1 2 3 4 5 6 7 8 9 10 11 12
BAC
(b)
Row
1
2
3
4

FIGURE 8

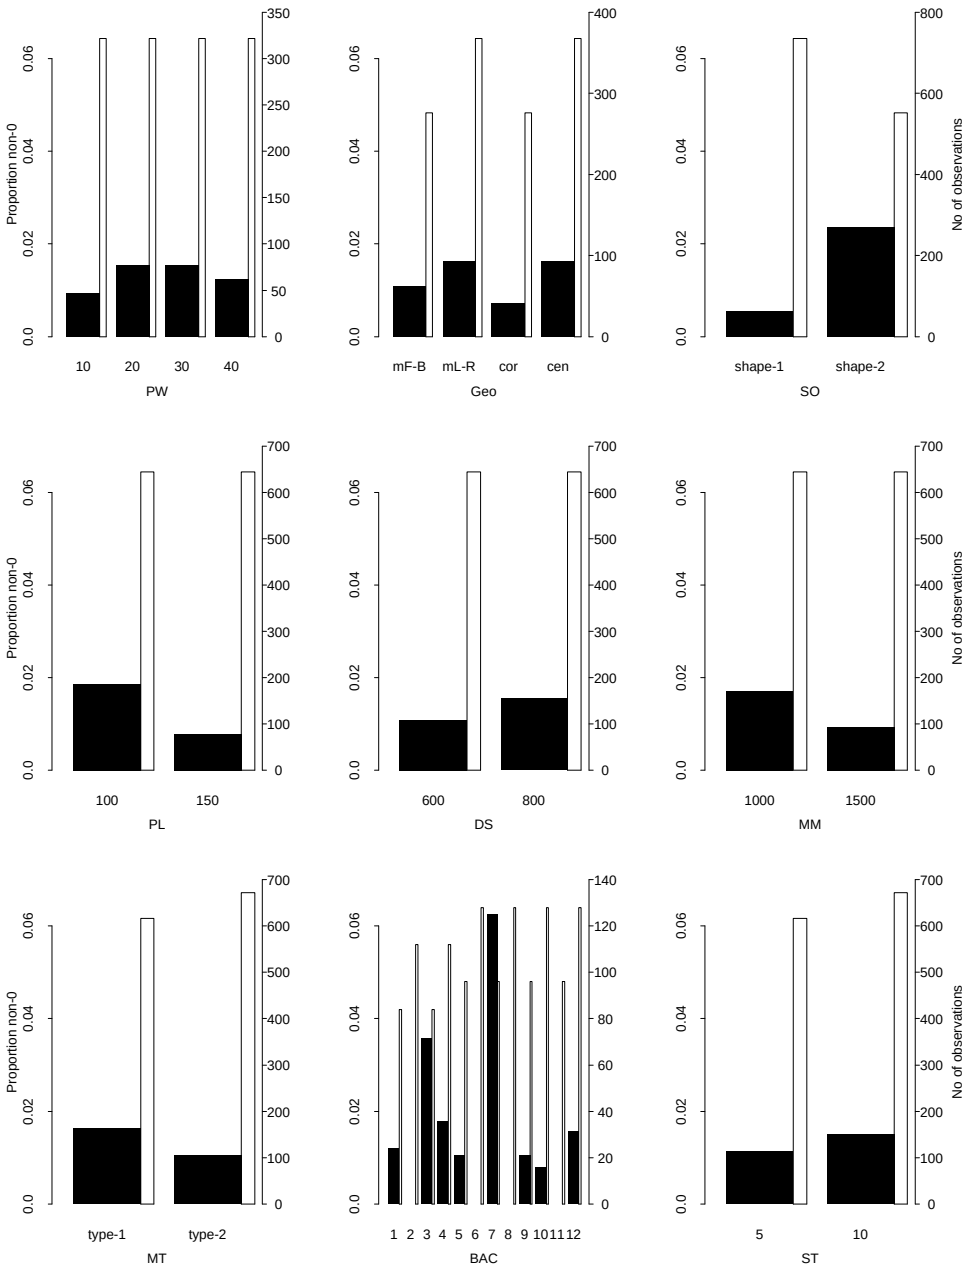
Proportion non-0
No of observations
PW
10
20
30
40
Geo
mF-B
mL-R
cor
cen
SO
shape-1
shape-2
PL
100
150
DS
600
800
MM
1000
1500
MT
type-1
type-2
BAC
1 2 3 4 5 6 7 8 9 10 11 12
ST
5
10

FIGURE 9