Appendix D – 10-Step Graphical Analysis Technique

Appendix D 10-Step Graphical Analysis Technique

We adopted a (NIST-developed) 10-step graphical analysis technique to evaluate the behavior of our model. The analysis technique uses 10 different plot types, listed in Table D-1, which allow us to identify the main factors influencing system behavior, to discover interactions among factors, to assess statistical significance of the factors, to propose linear models that match the data, to identify best and worse combinations of factors and to suggest additional factor settings that can drive system responses in particular directions. In this section, we introduce the technique through a sample sequence of the 10 plot types, using as an example one response variable, y11, average congestion window (CWND). The sample plots were taken from our study.

Table D-1. Identity and Purpose of 10 Plots in the 10-Step Graphical Analysis

Plot	Purpose
Ordered Data Plot (D.1)	Reveal how combinations of parameter settings influence response
Multi-factor Scatter Plot (D.2)	Reveal influence of individual parameter levels on response distribution
Main Effects Plot (D.3)	Reveal individual parameters having greatest influence on response
Interaction Effects Matrix (D.4)	Reveal degree of influence of parameter pairs on response
Block Plot (D.5)	Test robustness of statistically significant parameters in light of secondary or nuisance factors
Youden Plot (D.6)	Reveal parameters and parameter pairs with greatest influence on response
|Effects| Plot (D.7)	Reveal magnitude of a change in response due to specific parameters and parameter interactions
Half-Normal Probability Plot of |Effects| (D.8)	Separate influential parameters and parameter interactions from those that are not influential
Cumulative Residual SD Plot (D.9)	Provide information sufficient to construct a linear model to represent response data
Contour Plot (D.10)	Suggest how alterations in parameter settings could influence system response in predictable directions.

The transmission control protocol (TCP) manages a congestion window (CWND) variable that represents the number of packets that can be sent prior to receipt of an acknowledgment. The larger the CWND, the more packets that can be sent per unit time and thus the greater will be the transmission rate. For that reason, a network with a high average CWND (y11) will be able to transmit more packets than a network with a lower average CWND. In general, a CWND is reduced when packets are lost, usually due to congestion. Lowering the CWND slows the rate of packet transmissions in the network and thus should reduce congestion. Subsequent to a reduction, TCP increases the CWND linearly and so the rate of packet transmissions in the network should also increase. Once the transmission rate becomes too high, packets are lost and the CWND is reduced and the rate of transmission slows and so on. Thus the average CWND size might be used to represent the level of congestion in a network. Here, we analyze average CWND size with 10 plots, each representing one of the plot types listed in Table D-1. The changes in CWND were driven by a sensitivity analysis experiment, described in Chapter 4.

D.1 Ordered Data Plot. Fig. D-1 shows an ordered data plot, which graphs a system response (y axis) against every combination of factors (x axis) investigated in an experiment. The data is arrayed from smallest (on the left) to largest (on the right) response value. Our sensitivity analysis used 64 combinations of the 11 factors (recall Fig. 4-1) and so the plot contains 64 points. The upper left-hand corner of the plot shows the number of factors (k =11) and the number of combinations of factors (n = 64) in the experiment that generated the data. Below the x axis, each point is labeled with the specific combination of factor (x1 to x11) levels (- or +) that led to the response.

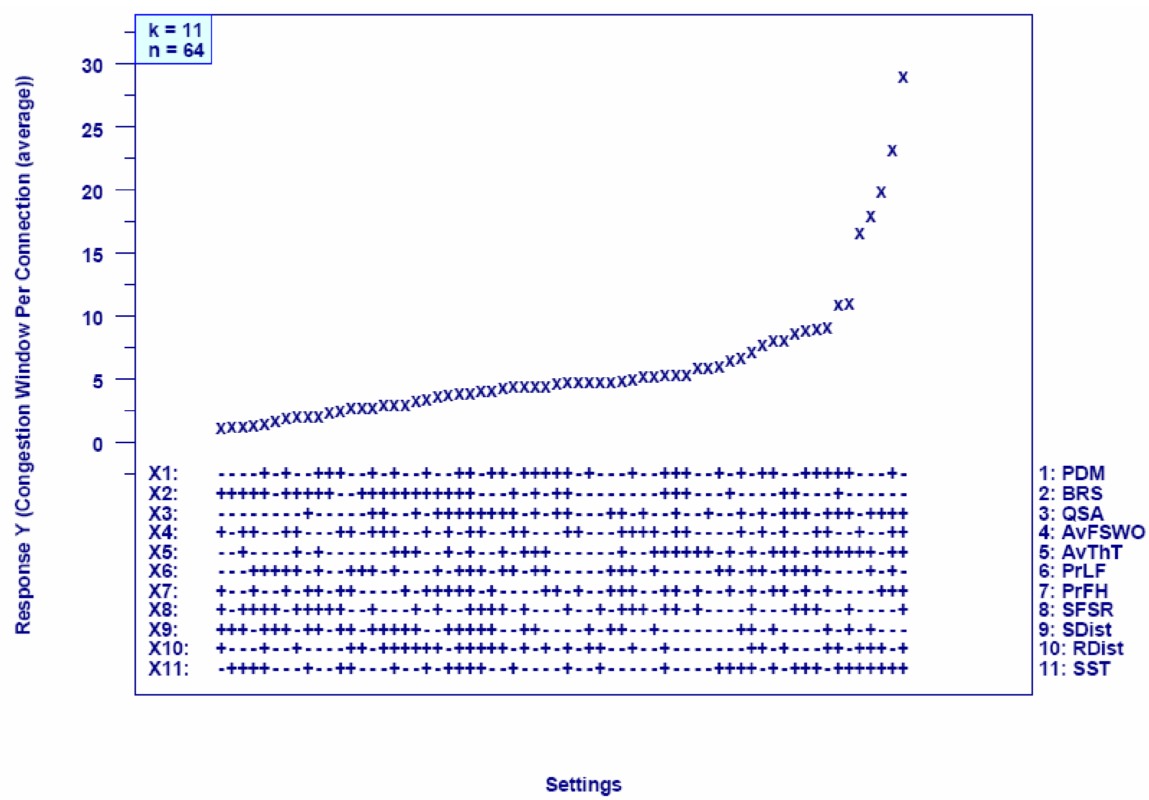

Figure D-1. Sample Ordered Data Plot

A legend to the right of the plot gives shorthand names to identify each of the 11 factors. Here the factors include: propagation delay (PDM/x1), network speed (BRS/x2), buffer-sizing algorithm (QSA/x3), average file size (AvFSWO/x4), average think time (AvThT/x5), probability that a user downloads a larger file (PrLF/x6), probability that a source resides on a fast host (PrFH/x7), scaling factor for the number of sources and receivers (SFSR/x8), distribution of sources (SDist/x9) and receivers (RDist/x10) throughout the topology, and initial slow-start threshold (SST/x11).

The plot can reveal the combination of factor settings that lead to the smallest (left-most) and largest (right-most) response from the system. In addition, the plot can reveal combinations of factor settings that appear to have greatest influence on the response. From Fig. D-1, we might conclude that the five right-most combinations had more significant influence (than other combinations) in increasing the average CWND per connection. Examining the factor settings associated with these data points reveals some common factors among them. For example, these points all have higher[1] network speed (x2 = -) and higher initial slow-start threshold (x11 = +) and four of the five points have larger (x3 = +) buffer sizes. Looking across the row of settings for network speed (factor x2) one can see that higher network speeds (x2 = -) seem to result more often in higher average CWND. Similar views can be taken of the other factors.

As a result of viewing the ordered data plot, an experimenter begins to see how various factors could be driving system response. From Fig. D-1 alone, an experimenter knowledgeable in the domain could begin to get a sense that faster networks with less congestion, higher initial slow-start threshold and larger buffer sizes lead to higher average CWND. More detailed information on the influence of these and other factors becomes available from subsequent plots.

D.2 Multi-factor Scatter Plot. Fig. D-2 shows a sample multi-factor scatter plot, which groups responses for each factor by setting (+ or -) and then plots the responses (y axis) together with the average response (dashed horizontal line). The plot also gives the number of factors (k = 11) and observations (n = 64). The x axis shows each factor (x1 to x11) as two vertical scatter plots, one when the factor setting is a minus and one when the factor setting is a plus. Thus, each individual scatter plot has half of the observations (here 32 of 64). The plot shows the distribution of response values and identifies the minimum and maximum values. From Fig. D-2 one can see that average CWND tends to be under 10 for most observations, which might suggest that many of the experiment combinations constrained CWND. The plot also reveals a clear setting for each factor in order for CWND to achieve its maximum value (near 30). This combination of factors will correspond with the right-most combination of factors in the ordered data plot.

Interpreting Fig. D-2, an experimenter can see the following settings leading to highest average CWND: shorter propagation delay (x1 = -), higher network speed (x2 = -), larger buffer sizes (x3 = +), larger file sizes (x4 = +), longer think times (x5 = +), higher probability of transferring larger files (x6 = -)[2], lower probability of fast hosts (x7 = +)[3], more sources (x8 = +), less uniform distribution of sources (x9 = -), more uniform distribution of receivers (x10 = +), and higher initial slow-start threshold (x11 = +). The experimenter might wonder which of the factors and settings are most influential. The next plot provides this information.

---

[1] Recall the miscoding of factor x2: minus is higher network speed and plus is lower network speed.
[2] The coding for factor x6 was reversed (plus was a lower probability and minus was a higher probability).
[3] The coding for factor x7 was also reversed (plus was a lower probability and minus was a higher probability).

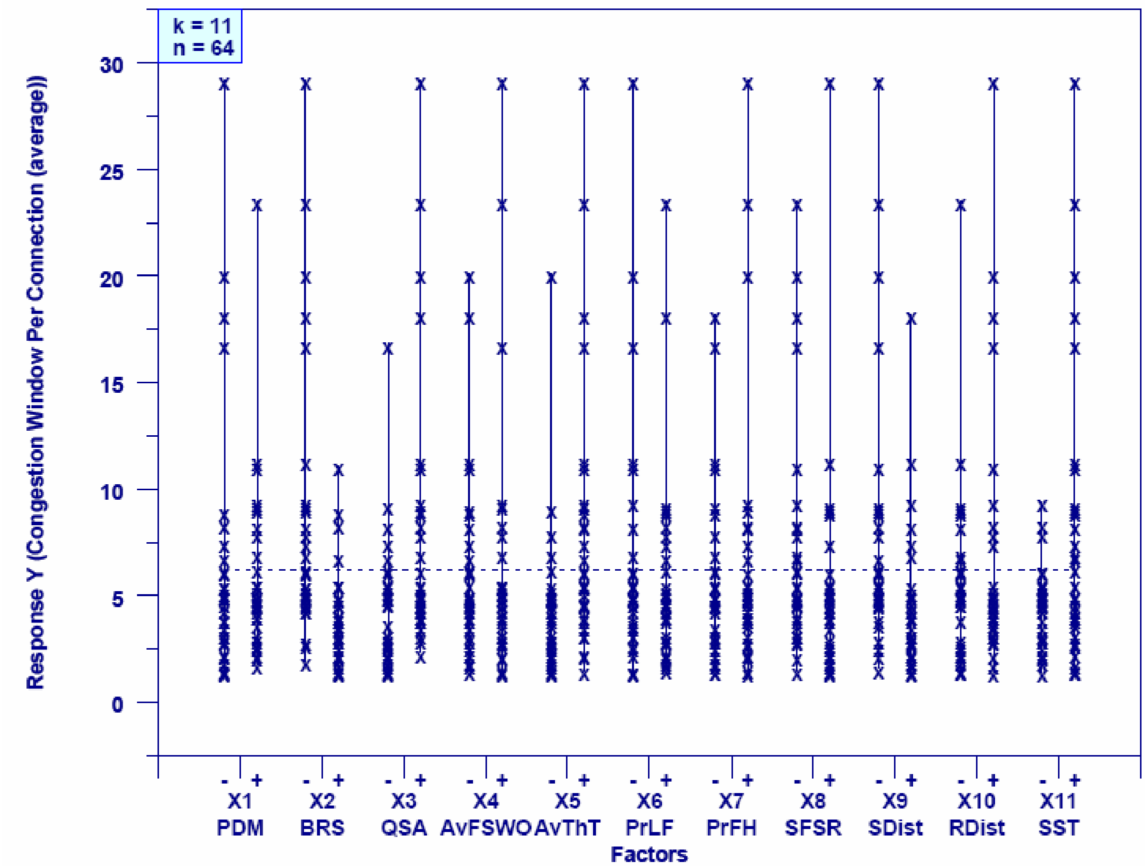

Figure D-2. Sample Multi-factor Scatter Plot

D.3 Main Effects Plot. Fig. D-3 gives a sample main effects plot, which is the most essential plot to identify the factors and settings driving a system's response. The basic framing of the plot is similar to that of the multi-factor scatter plot; however, each vertical scatter plot is replaced by an average of the response. For each factor, the averages are connected with a line that indicates the magnitude and direction the response changes when moving from a minus to a plus setting for the factor. On the x axis, each factor is annotated with the absolute change in response and the change relative to (i.e., as a % of) the mean response.

Fig. D-3 reveals that the most influential factor in determining CWND is network speed (70 % of the mean) followed by three closely grouped factors: buffer-sizing algorithm (54 %), initial slow-start threshold and think time (each 53 %). The distribution of sources also has a significant (50 %) influence. Notice that the plot reveals a smaller number of sources and receivers (x8 = -) leading to a (1.7 packet) larger average CWND than a larger number of sources – this is true despite the fact that the ordered data plot and scatter plot showed that the largest CWND was achieved when the number of sources was at its higher setting. In fact, a domain expert will understand that fewer sources sharing the same network means that each source may transmit faster, which is reflected in a larger CWND. Thus, here the main effects plot clearly reveals the true nature of the influence of the factors and settings on the response.

In thinking about the main effects, an experimenter with domain knowledge might be quite pleased with the meaning of these results regarding the validity of the model. Fewer, simultaneously active, flows (x5 = +, x8 = - and x9 = -), higher network speeds (x2 = -) together with more buffers (x3 = +) should permit higher CWND. Under these circumstances, the ability to increase the CWND to a higher threshold via initial slow-start (x11 = +) should also lead to higher CWND, because CWND increases fastest during initial slow start.

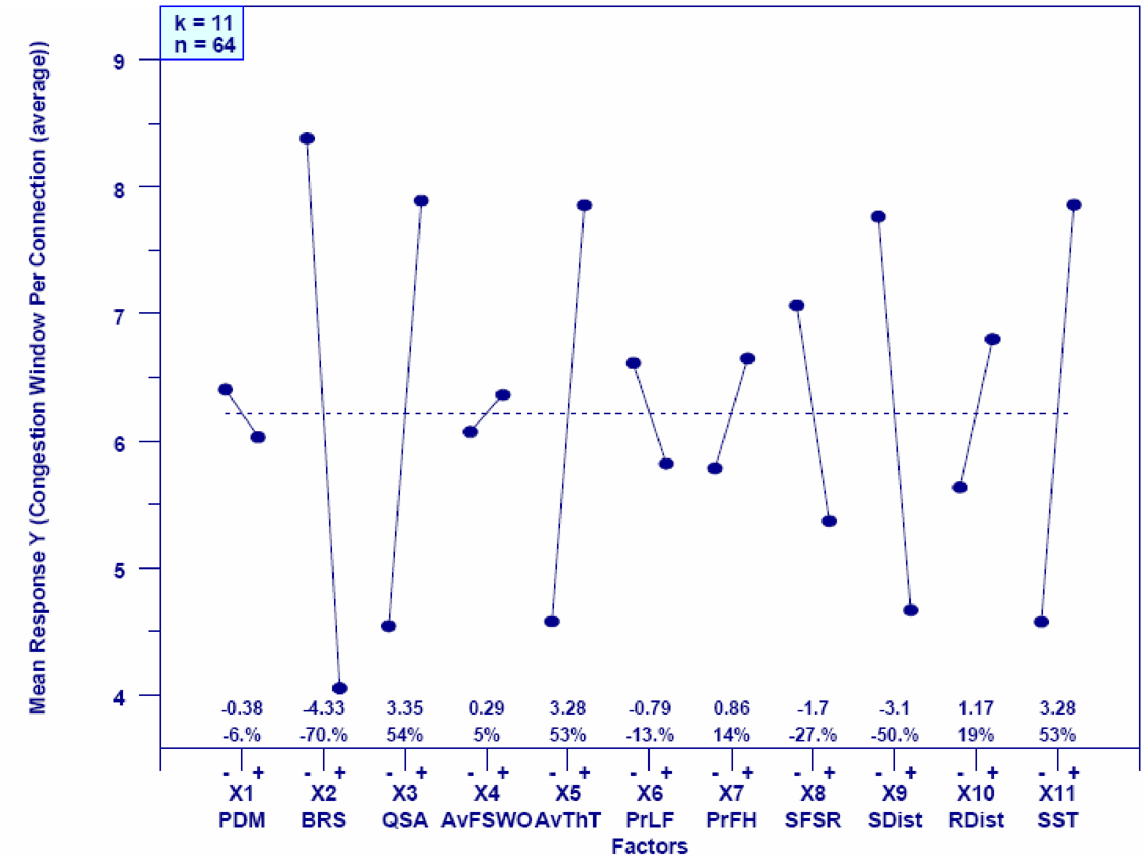

Figure D-3. Sample Main Effects Plot

D.4 Interaction Effects Matrix. Figure D-4 shows a sample interaction effects matrix. The purpose of this plot is to determine if interactions among factors have a significant influence on the response. If this plot reveals no such interaction effects, then an experimenter can conclude that the system response is driven primarily by main effects. The plot might also reveal interactions that an experimenter expected based on domain knowledge. On the other hand, the plot could reveal significant, unexpected effects due to interactions, which requires further investigation by the experimenter.

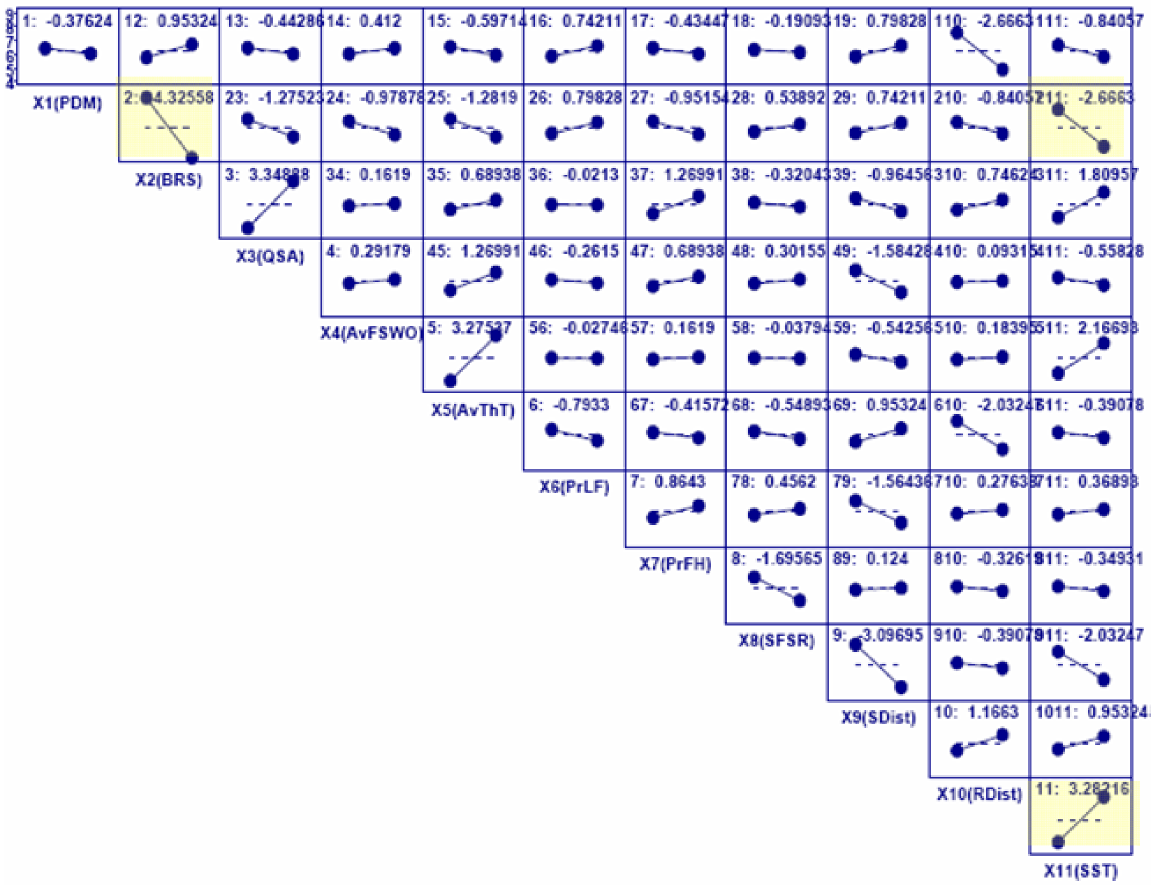

Figure D-4. Sample Interaction Effects Matrix

The interaction effects matrix takes the form of a half matrix containing rows and columns of sub-plots, where each sub-plot shows how the average response changes when moving from a minus to a plus setting for some combination of factors. The left-most sub-plot in each row (also the bottom-most sub-plot in each column) gives the main effects plot (from Fig. D-3) for a specific factor (x1 in the top row and x11 in the bottom row). Each of the remaining sub-plots in a given row (or column) show how the average response changes when moving from a minus to a plus setting for two factors (the factor beginning the row or column and each of the other factors).

An experimenter may scan the matrix starting from each main effects plot. Scan up (the related column) and also right across (the related row) to compare the influence of each main factor to the influence of possible two-factor interactions. Scanning the matrix in Fig. D-4 shows that, for the most significant main effects (x2, x3, x5, x8, x9 and x11), the influence of the main effect is greater than the influence of any interactions. So, for example, consider the three sub-plots highlighted in Fig. D-4. The upper left-hand sub-plot reports that changing network speed (x2) changes CWND size by 4.3 packets and the lower right-hand sub-plot reports that changing initial slow-start threshold (x11) changes CWND size by 3.3 packets, while the third highlighted sub-plot reports that changing x2 and x11 together changes CWND size by only 2.7 packets. Similar results can be found when comparing the influence of other factors with their two-factor interactions. This suggests that system response is driven by main effects and not two-factor interactions.

D.5 Block Plots. Block plots provide an elementary test of statistical significance for the influence of main effects. Given a full factorial experiment design, one can compare the average response for a minus setting of a factor to a plus setting under all possible combinations of other factors. This can provide a large amount of visual information, so typically only a subset of these plots is generated. Further, given an orthogonal fractional factorial (OFF) experiment design a reduced amount of information is available for generating block plots, so such plots are not as useful for OFF experiment designs. Still, block plots can prove useful in confirming findings about main effects.

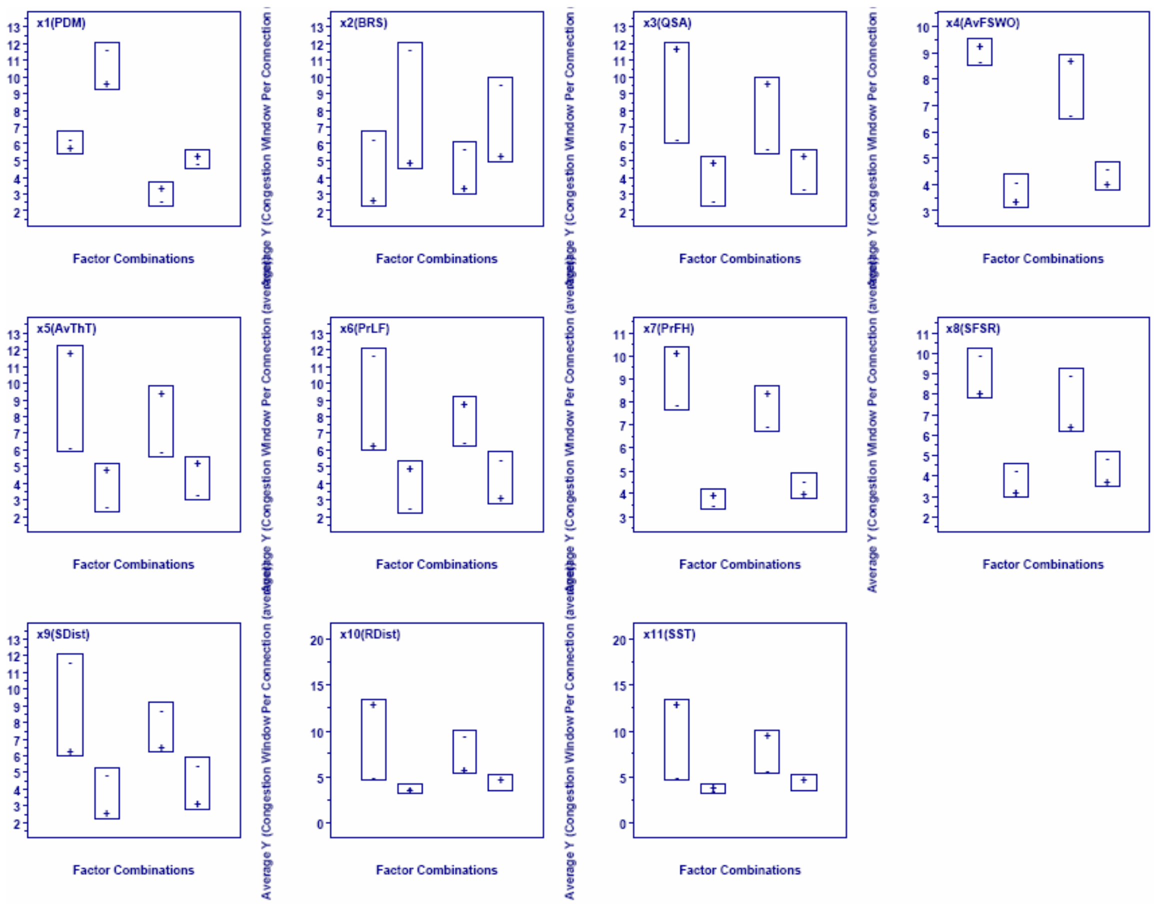

Figure D-5. Sample Block Plots

Fig. D-5 shows sample block plots for each of the 11 factors used in the sensitivity analysis of MesoNet. Each plot shows the average response for four combinations of secondary factors when the main factor of the plot is set to a minus and a plus. The block plots reinforce the findings of the main effects plot: the most significant factors are network speed (x2), buffer sizes (x3), think time (x5), number (x8) and distribution (x9) of sources and initial slow-start threshold (x11). This is revealed by the fact that one particular setting for each of these factors always leads to a higher value for

CWND. This is true no matter what combination of other factors is used. Results are mixed for the other five factors.

D.6 Youden Plot. A Youden plot, see Fig. D-6, graphs the average response for each factor (and two-factor interaction) when the factor (or factors) are set to a minus against the average response when set to a plus. In Fig. D-6 the average CWND for minus settings are plotted on the x axis and for plus settings on the y axis. For unimportant factors, the values should be nearly the same (and appear in the center of the graph). For important factors, values should lie toward the upper-left and lower-right corners of the graph.

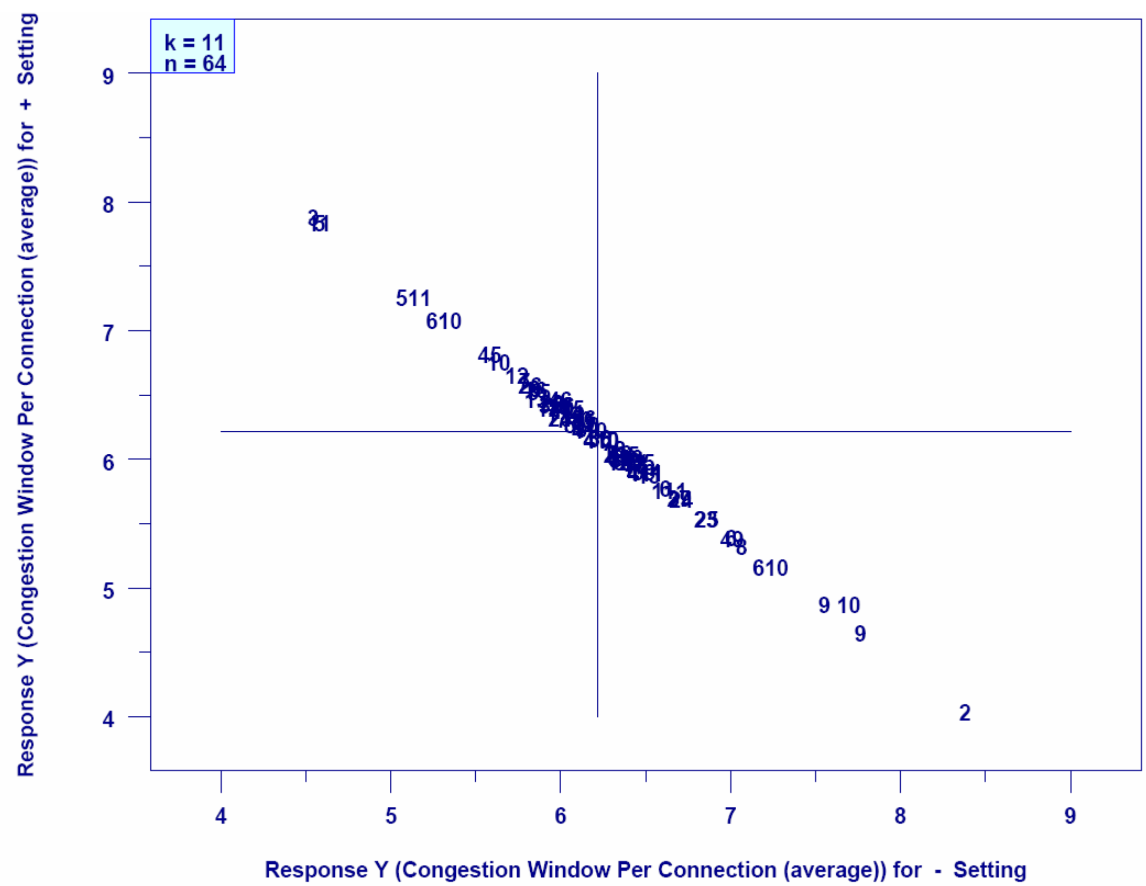

Figure D-6. Sample Youden Plot

Examining Fig. D-6 reveals that the most important factors are network speed (x2), buffer size (x3), think time (x5), source distribution (x9) and initial slow-start threshold (x11). The number of sources (x8) is not as important. Recall that the main effects plot identified x8 as of less importance than the other effects. The Youden plot supports the earlier finding. The plot also reveals some information about the influence of interactions. The distribution of receivers (x10) has a combined effect with the distribution of sources (x9) and the think time (x5) has a combined effect with initial slow-start threshold (x11). These interaction effects (also revealed in the interaction effects matrix, Fig. D-4) are less important than the main effects.

D.7 |Effects| Plot. The |Effects| plot, Fig. D-7, displays the absolute magnitude of a change in response due to specific factors and interactions. The x axis identifies the factors or interactions for which the corresponding magnitude of the change in response is plotted on the y axis. The factors are ordered by decreasing magnitude from left-to-right on the x axis. This plot should confirm the information given in previous plots regarding the influence of factors and interactions. The average value for the response is given in the upper left-hand corner of the plot. The plot may be augmented (as here) with a rank-ordered list of factors (and interactions) and the associated (signed) magnitude of the effects.

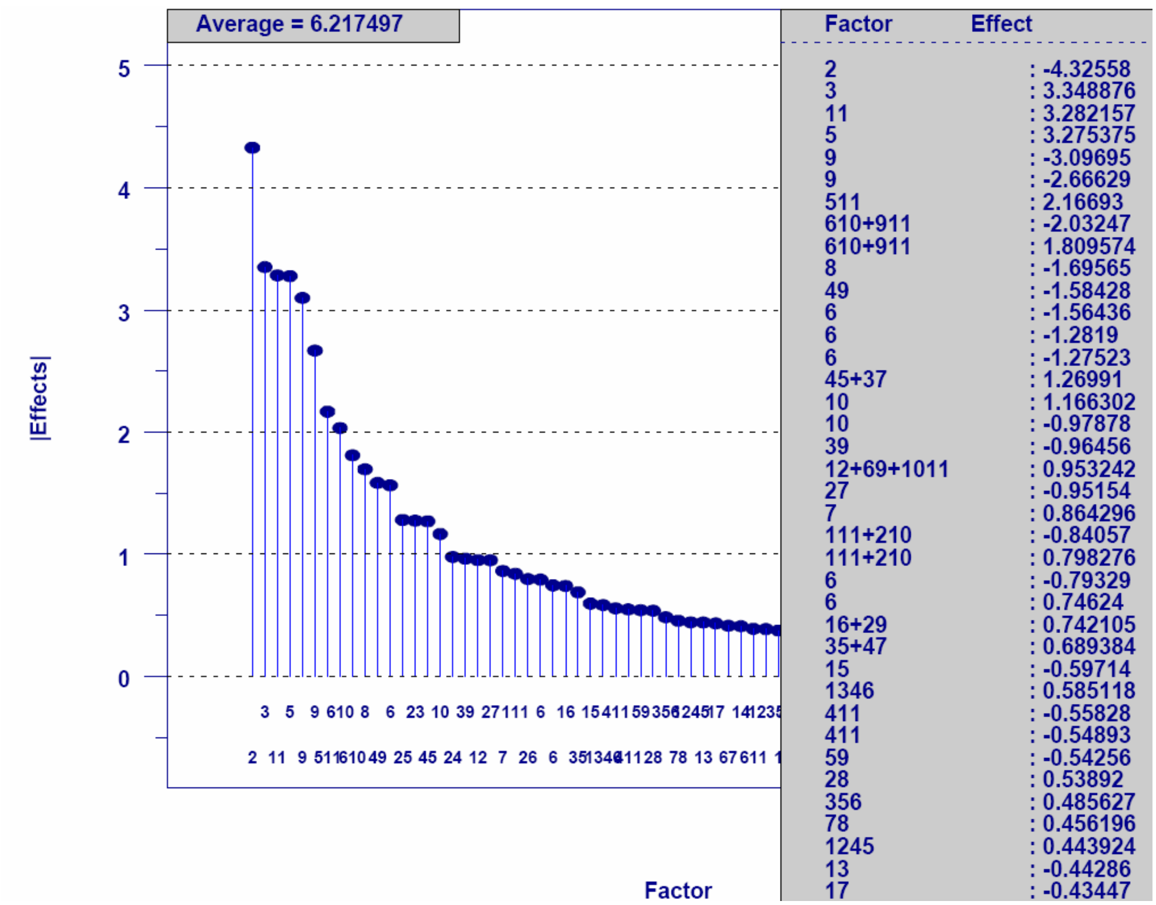

Figure D-7. Sample |Effects| Plot

Fig. D-7 confirms earlier findings: the main factors influencing CWND include (in order) network speed (x2), buffer size (x3), initial slow-start threshold (x11), think time (x5) and source distribution (x9). The associated list identifies x9 twice, but this is a mistake in labeling. Consulting the interaction effects matrix (Fig. D-4) reveals that the second x9 should be an x2-x11 interaction. This demonstrates the cross-checking value of the redundancy included in the 10-step graphical analysis technique.

D.8 Half-Normal Probability Plot of |Effects|. A half-normal probability plot of |Effects|, Fig. D-8, classifies effects as important or unimportant. The x axis of a half-normal probability plot represents the ordering of a theoretical half-normal distribution of values. The y axis represents the |Effects|, which are plotted in order from least to greatest. Unimportant effects would tend to have mean difference centered on zero, sp when plotting the |Effects| for such data, one would expect plotted values to begin around zero and ascend linearly. Linear data emanating from the origin implies that the effects have a non-significant (zero) value. Data departing from linearity indicate a statistically significant effect.

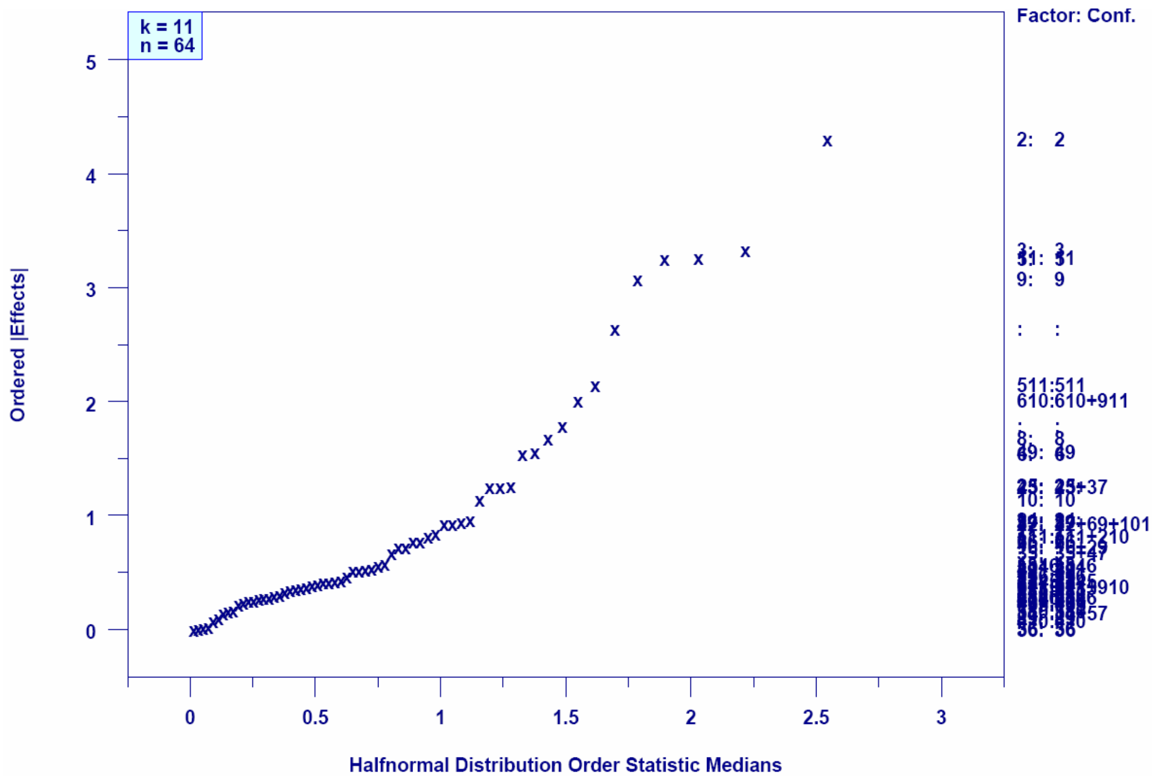

Figure D-8. Sample Half-Normal Probability Plot of |Effects|

As illustrated in Fig. D-8 the values to the right of the plot identify which factor (or interaction) is responsible for the plotted value on the y axis. Interpreting Fig. D-8 indicates the value of CWND is driven mainly by five factors: network speed (x2), buffer size (x3), think time (x5), initial slow-start threshold (x11) and source distribution (x9). The plot shows relative importance: x2, followed by the grouping of x3, x5, x11 and then finally x9. This finding is consistent with information obtained from previous plots.

D.9 Cumulative Residual Standard Deviation Plot. The cumulative residual standard deviation (SD) plot (Fig. D-9) provides information sufficient to construct a linear model to represent experiment data. The y axis of the plot gives the residual error between the data and a fitting function when adding terms representing the set of factors and interactions represented on the x axis. The first term on the x axis is the residual SD when describing the data using only the grand average over all factors. Then, the residual standard deviations are plotted as the influence of each factor is added in order of decreasing reduction in SD. The factors are identified on the x axis.

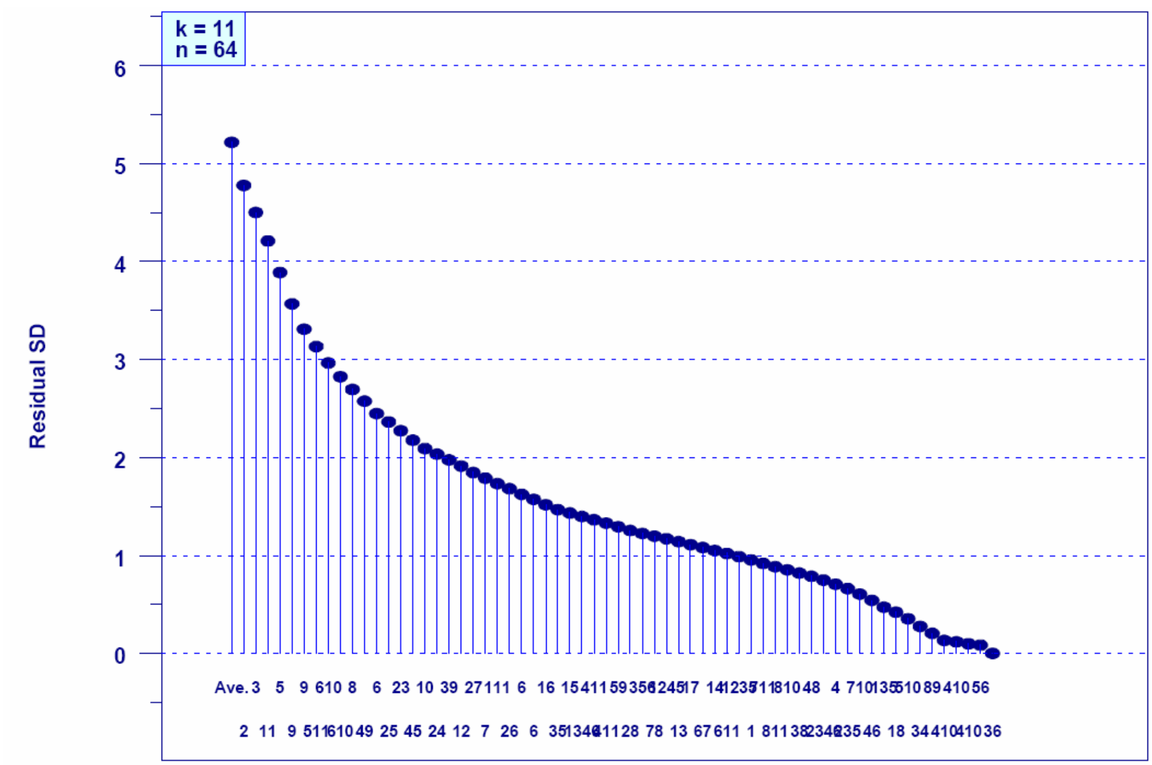

Figure D-9. Sample Cumulative Residual Standard Deviation Plot

Typically, one hopes for a model where the most important factors can explain most of the standard deviation in the data. If this holds, then a fairly simple linear model can describe the data. A plot demonstrating such a case would exhibit a large reduction in SD after the main factors are included in the model. In our example, when we include the five most important factors (x2, x3, x5, x9 and x11), Fig. D-9 shows a remaining error of about 3.5. The encouraging news from Fig. D-9 is that the five most important factors reduce the SD the most (and in the expected order). On the other hand, the residual SD is still rather high after including these factors, so the resulting linear model could likely not be used for interpolation or prediction.

D.10 Contour Plot of Two Dominant Factors. The contour plot aims to suggest other factor settings that could alter system response in predictable directions. While one could plot each pair of factors together, a typical approach is to construct a contour plot from the two most important factors. Fig. D-10 shows a sample contour plot from the two main factors, x2 (network speed) and x3 (buffer size), influencing CWND. The most important factor is plotted on the x axis and the second most important factor is plotted on the y axis. The axes are labeled with an origin (0) and then the two settings (-1 and +1) for each factor. A point is placed at each combination of factors (x2, x3) = {(-1, -1), (+1, -1), (+1, +1), (-1, +1)}. These four points are connected with a dashed line to form a rectangle. Each point is labeled with the value of one of the specific response variables for the associated combination of the two factors. Based on the fitted model developed when generating the cumulative residual SD plot, contour lines are added to form a contour plot.

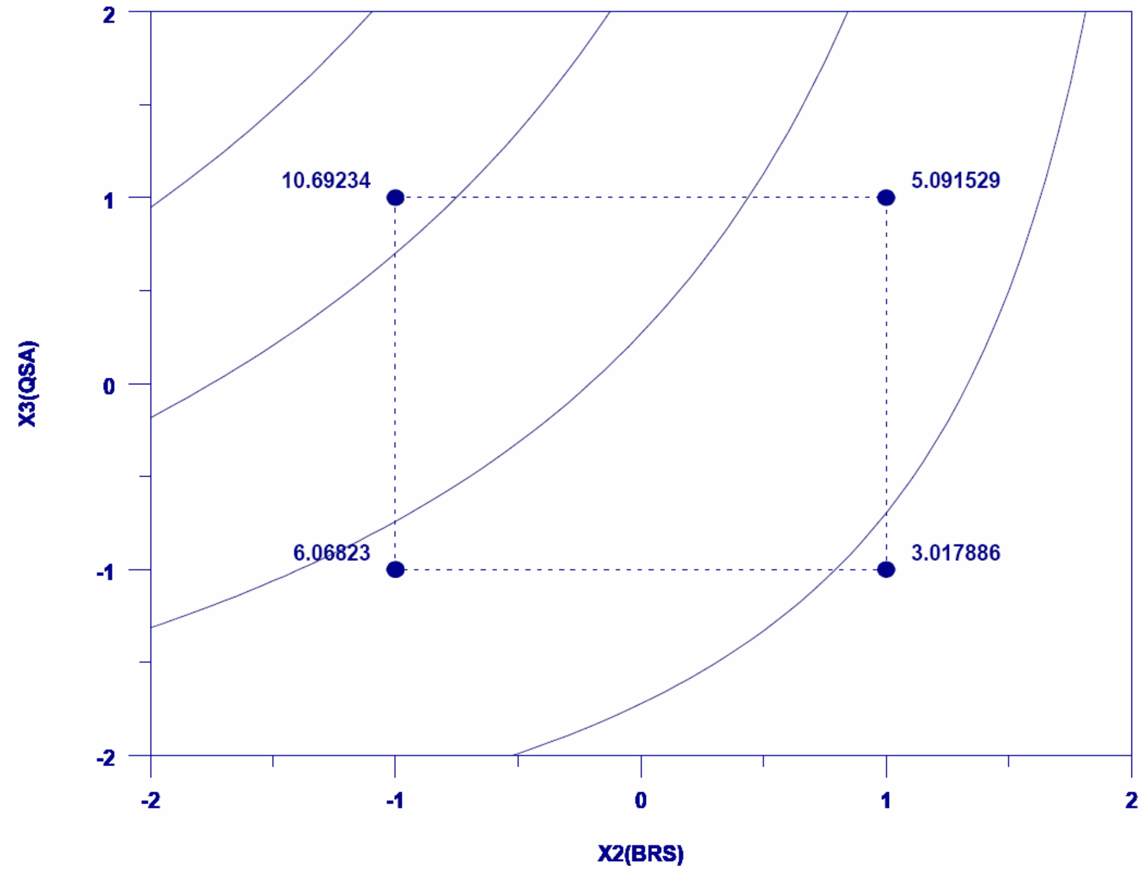

Figure D-10. Sample Contour Plot of Two Dominant Factors

As Fig. D-10 shows, the combination of (-1, +1) – higher network speed and larger buffer size – produces the largest CWND (10.69…). The contour lines indicate that increasing network speed and buffer size would lead to larger CWND values, while decreasing network speed and buffer size would lead to smaller CWND values. Under some experiments and conditions, the contour plot could reveal directions in which to alter factor settings to create more optimal responses or suggest how responses might change as factor settings are altered.