How To: Analyze a Split-Plot Design

Using STATGRAPHICS Centurion

by

Dr. Neil W. Polhemus

August 13, 2005

Introduction

When performing an experiment involving several factors, it is best to randomize the order in which the experiments are performed. This reduces the chance that any unexpected effects, such as a gradual change in some uncontrolled variable over time, will bias the results of the experiment. Randomization protects the experimenter against the effects of "lurking variables", which are factors that affect the experiment but are not recognized until after the experiment is completed (if they ever are). Randomization also insures that all experimental factors are subjected to the same level of experimental error, which simplifies the type of analysis that needs to be performed on the results.

In some cases, pure randomization is not practical, since certain factors may be hard to change. For example, when experimenting with a chemical process, it may be very easy to change the initial temperature by turning a dial, but changing a factor such as the type of catalyst may be considerably more involved. In such cases, experimental designs such as the split-plot design provide economical alternatives to full randomization. Unfortunately, the restricted randomization complicates the statistical analysis that must be performed on the resulting data.

This "How To" guide shows how STATGRAPHICS Centurion can be used to analyze typical split-plot designs. Two examples are considered, one involving categorical experimental factors and the other involving quantitative factors.

Example #1

The first example we will consider comes from the latest edition of Statistics for Experimenters, second edition (Wiley, 2005) by Box, Hunter and Hunter. It is an experiment designed to study the corrosion resistance of steel bars that have been treated with four different coatings. The bars have been randomly positioned in a furnace and baked at three different temperatures. Although the position of the bars in the furnace could be randomized, multiple experiments involving a particular temperature needed to be conducted at the same time, since it was impractical to change the temperature of the furnace for each sample.

The following table shows the layout of the experiment. Six experimental runs were performed, each at a selected temperature. During each run, four bars were baked, one bar with each coating. The position of the bars in the oven was randomly determined for each run. However, the temperature was changed in a systematic manner:

Run	*Temperature*	*Coating of Bar by Position in Furnace*
1	360	C2, C3, C1, C4
2	370	C1, C3, C4, C2
3	380	C3, C1, C2, C4
4	380	C4, C3, C2, C1
5	370	C4, C1, C3, C2
6	360	C1, C4, C2, C3

Figure 1: Experimental Design

The results of the experiment have been placed in a file called *howto11.sf6*, which has the structure shown below:

Temperature	*Replicate*	*Coating*	*Corrosion*
360	1	2	73
360	1	3	83
360	1	1	67
360	1	4	89
370	1	1	65
370	1	3	87
370	1	4	86
370	1	2	91
380	1	3	147
380	1	1	155
380	1	2	127
380	1	4	212
380	2	4	153
380	2	3	90
380	2	2	100
380	2	1	108
370	2	3	150
370	2	1	140
370	2	3	121
370	2	2	142
360	2	1	33
360	2	4	54
360	2	2	8
360	2	3	46

Figure 2: Experimental Results in File Howto11.sf6

The goal of the experiment is to determine the effect of the type of coating and the temperature on corrosion.

Step 1: Plot the Data

The first step when analyzing any new data set is to plot it. In this case, a coded scatterplot is very useful.

Procedure: X-Y Scatterplot

To plot the experimental data, let's begin by pushing the *X-Y Scatterplot* button 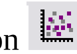 on the main toolbar. On the data input dialog box, indicate the variables to be plotted on each axis as shown below:

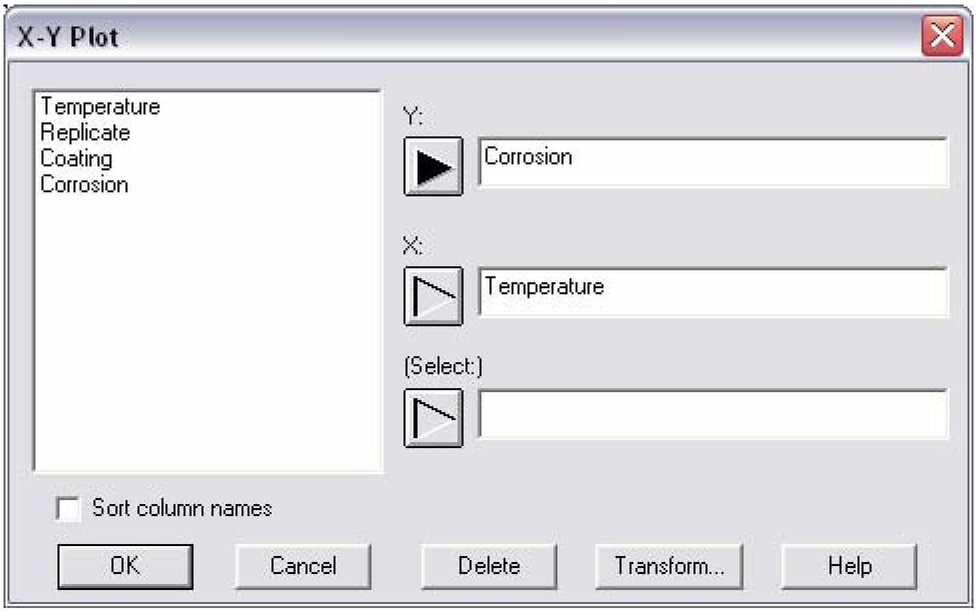

Figure 3: Data Input Dialog Box for X-Y Scatterplot

The resulting plot shows a general increase in corrosion with increasing temperature:

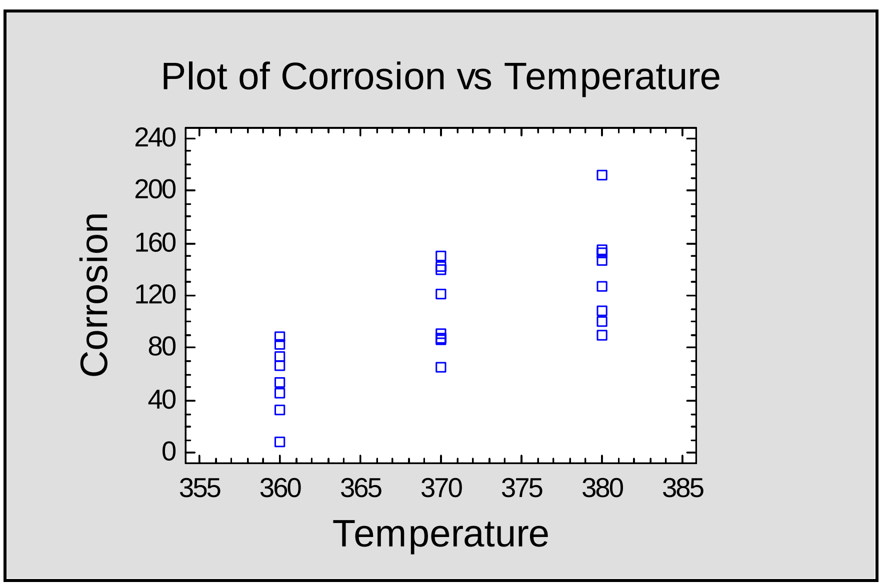

Figure 4: X-Y Scatterplot for Chlorine Data

To introduce the type of coating into the plot, double-click on the graph to enlarge it and press the *Pane Options* button 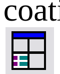 on the analysis toolbar. This will display the following dialog box:

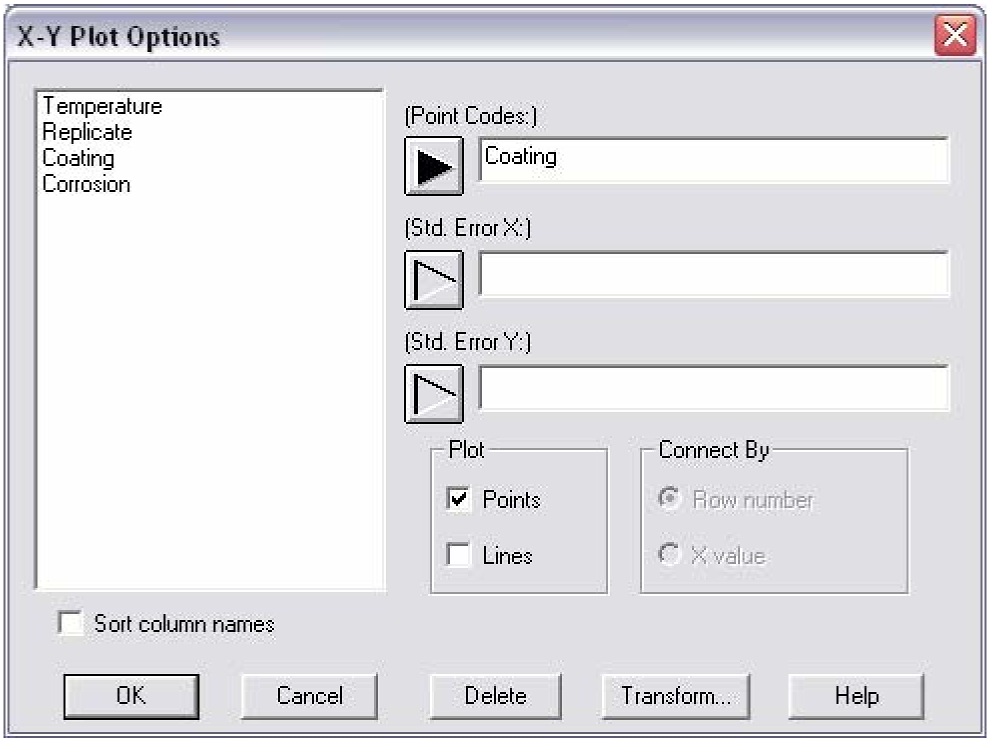

Figure 5: X-Y Plot Options Dialog Box

Enter *Coating* in the *Point Codes* field to generate a coded scatterplot:

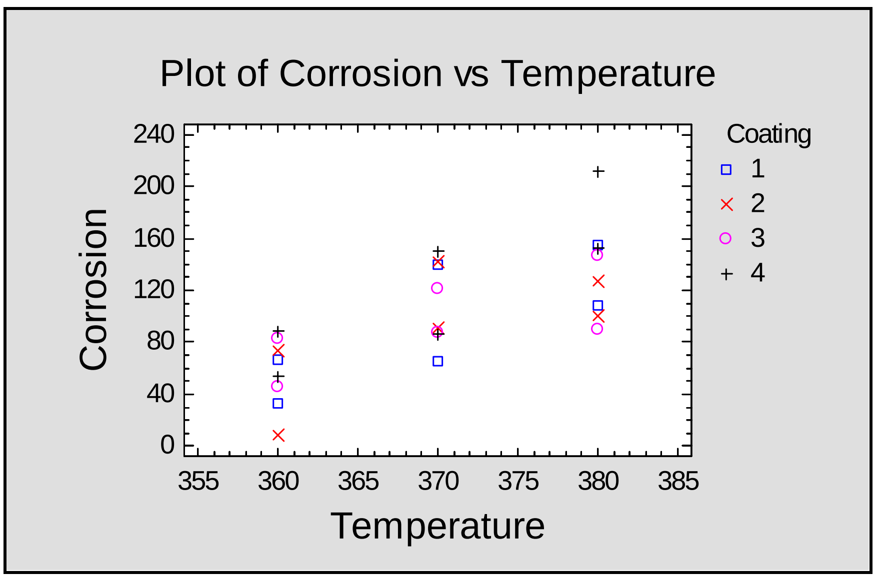

Figure 6: Coded X-Y Scatterplot

Each type of point symbol represents a different coating. At each of the 3 temperatures, there are 2 replicates for each type of coating. Based on this plot, one would be hard pressed to select the best type of coating.

Step 2: Analyze the Data

The corrosion experiment was conducted to determine the effect of two factors, *temperature* and type of *coating*, on a single response variable, *corrosion*. As in a standard factorial experiment, data has been collected at all combinations of the levels of the two experimental factors. What makes this experiment special is that the size of the experimental unit for one factor is different than for the other, causing the experimental error affecting the estimates of the factor effects to be different. In particular, the furnace temperature was set only six times, corresponding to the six rows in Figure 1. Each of these rows is called a “whole-plot”, with analogies to early agricultural experiments in which one type of fertilizer might be applied to an entire field. Within each whole-plot, the coatings are randomly allocated to positions within the furnace. Position is referred to as the “subplot”.

There are two sources of variability in this experiment: variability in bringing the furnace to a particular temperature, and variability amongst positions within the furnace. The variance of the first source will be labeled σ_W^2, while the variance of the second will be labeled σ_S^2. Since coatings were randomized across the subplots, estimates of differences between the coatings are subject only to the subplot error, while estimates of the temperature effects involve both the subplot and whole-plot error.

The secret to analyzing this split-plot experiment is to view it as two experiments, one contained within the other. The whole-plot experiment involves changes in temperature. It can be viewed as a one-way comparison amongst the three levels of temperatures, with two replicates at each level:

T_1: 360ºC	T_2: 370ºC	T_3: 380ºC
R_1,R_2	R_1,R_2	R_1,R_2

Each replicate represents a “whole-plot” or run of the furnace. A standard ANOVA table for the whole-plot design would look like:

Source of Variation	*Degrees of Freedom*
Temperature	2
Error (Replicates)	3

The second experiment crosses coating with the whole-plots:

	T_1 R_1	T_2 R_1	T_3 R_1	T_1 R_2	T_2 R_2	T_3 R_2
C_1	Y	Y	Y	Y	Y	Y
C_2	Y	Y	Y	Y	Y	Y
C_3	Y	Y	Y	Y	Y	Y
C_4	Y	Y	Y	Y	Y	Y

One bar was baked at each of the 4×6 combinations of coating and whole-plot. The terms in the ANOVA table for this experiment look like:

Source of Variation	*Degrees of Freedom*
Whole-Plots	5
Coating	3
Temperature×Coating	6
Error	9

Tests of significance for the terms in each experiment are done by comparing them to the corresponding error component.

Procedure: General Linear Models

A fully randomized factorial design, in which there is only one source of experimental error, is usually analyzed using the *Multifactor ANOVA* procedure. When there is more than one source of error, the *General Linear Models* procedure must be used instead. This is accessed from the main STATGRAPHICS Centurion menu by selecting:

- If using the Classic menu: *Compare – Analysis of Variance – General Linear Models.*
- If using the Six Sigma menu: *Improve – Analysis of Variance – General Linear Models.*

The data input dialog box is shown below:

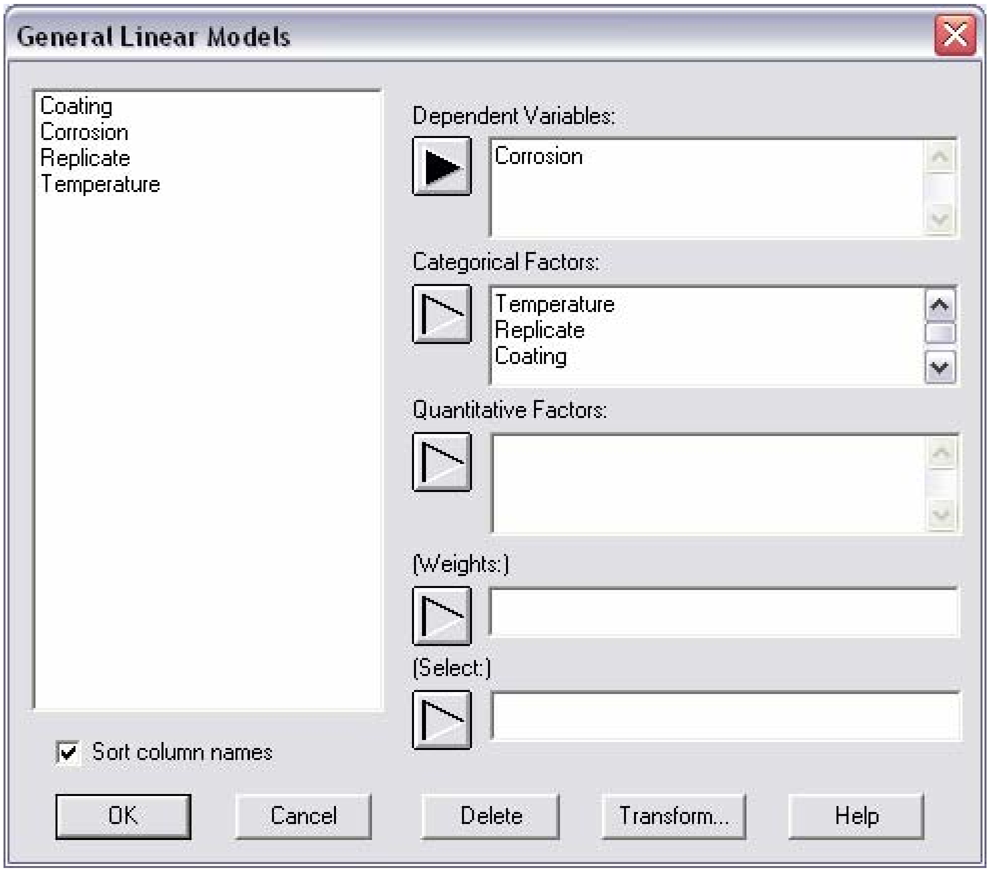

Figure 7: Data Input Dialog Box for General Linear Models

Although not strictly necessary, it is helpful to enter the factors in the order shown:

1. Whole-plot factors before subplot factors.
2. *Replicate* after *Temperature*, since the replicates are performed at each temperature level.

After completing the first dialog box, a second dialog box is displayed on which to specify the statistical model. It should be completed as shown below:

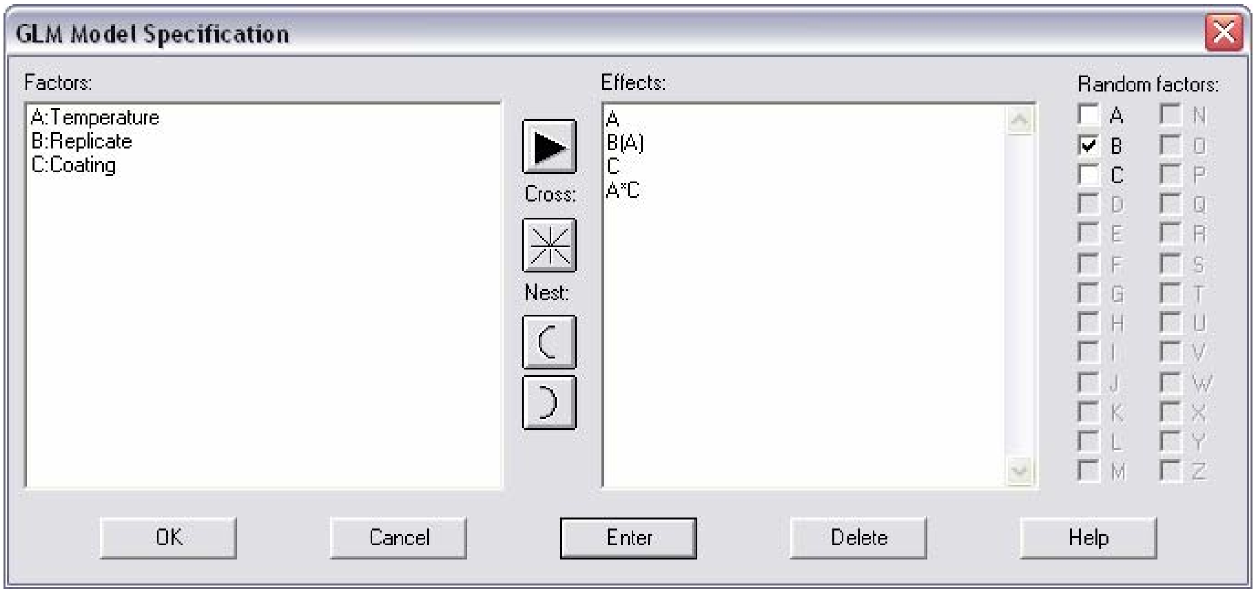

Figure 8: Model Specification Dialog Box for General Linear Models

Note the following:

1. The main effect of the whole-plot factor *Temperature* is specified by placing the single letter A in the *Effects* field. Since specific values of temperature were used (360, 370, 380), it is a fixed rather than a random factor.

2. *Replicate* is entered using the notation B(A). This indicates that replicate (Factor B) is nested within temperature (Factor A). Specifying the factors as nested indicates that the experimental unit for the first replicate of T_1 is not the same as the experimental unit for the first replicate of temperature T_2 or T_3. Indeed, each combination of T and R forms a separate whole-plot. Factor B is also specified to be a *random* factor, since the two replicates are but a small random sample of all replicates that could have been performed.

3. After entering the whole-plot terms, the subplot effects are specified. First, main effects of *Coating* are requested using the symbol *C*. Then the *Temperature*×*Coating* interaction is entered as *A**C*. Since four specific coatings are being tested, *Coating* is also a fixed rather than a random factor.

Pressing OK causes the specified model to be fit. The *Analysis Summary* pane summarizes the fitted model. The top section of that summary is shown below:

General Linear Models

Number of dependent variables: 1
Number of categorical factors: 3
A=Temperature
B=Replicate
C=Coating
Number of quantitative factors: 0

Analysis of Variance for Corrosion

Source	*Sum of Squares*	*Df*	*Mean Square*	*F-Ratio*	*P-Value*
Model	48517.8	14	3465.55	27.83	0.0000
Residual	1120.88	9	124.542		
Total (Corr.)	49638.6	23			

Type III Sums of Squares

Source	*Sum of Squares*	*Df*	*Mean Square*	*F-Ratio*	*P-Value*
Temperature	26519.3	2	13259.6	2.75	0.2093
Replicate(Temperature)	14439.6	3	4813.21	38.65	0.0000
Coating	4289.13	3	1429.71	11.48	0.0020
Temperature*Coating	3269.75	6	544.958	4.38	0.0241
Residual	1120.88	9	124.542		
Total (corrected)	49638.6	23			

Figure 9: GLM Analysis Summary – Top Section

The most important information in the above table is in the section labeled *Type III Sums of Squares*. The rightmost column of that table contains a P-Value for each term in the model. P-Values less than 0.05 indicate effects that are statistically significant at the 5% significance level. Both the main effect of *Coating* and the *Temperature*×*Coating* interaction are significant.

Also included in the *Analysis Summary* is the table shown below:

F-Test Denominators

Source	*Df*	*Mean Square*	*Denominator*
Temperature	3.00	4813.21	(2)
Replicate(Temperature)	9.00	124.542	(5)
Coating	9.00	124.542	(5)
Temperature*Coating	9.00	124.542	(5)

Variance Components

Source	*Estimate*
Replicate(Temperature)	1172.17
Residual	124.542

Figure 10: GLM Analysis Summary – Bottom Section

The *F-Test denominators* indicate which line in the ANOVA table has been used to test the significance of each effect. For *Temperature*, the (2) indicates that it has been compared against the whole-plot error on the second line of the ANOVA table (labeled *Replicate(Temperature)*). The other factors have been compared to the *Residual* or subplot error in line 5.

Also shown are the estimates of the error components:

Whole-plot error variance: $\hat{\sigma}_W^2 = 1172.17$

Subplot error variance: $\hat{\sigma}_S^2 = 124.542$

Notice that the subplot error, which represents differences between positions in the furnace, is much smaller than the whole-plot error, which represents differences between runs of the furnace. From the discussion in Box, Hunter and Hunter, it seems that this is evidence of difficulties in maintaining the desired temperature when an experiment is run.

Step 3: Display the Results

Once the important factors have been identified, it is useful to display the estimated effects graphically. To plot the main effects of each factor, press the *Graphs* button on the analysis toolbar. If you select *Means Plot*, the mean response will be plotted at each level of a selected factor. Using *Pane Options*, you can also choose an uncertainty interval to place around each mean. In the plot below, the mean value of *Corrosion* is shown for each of the three temperatures, together with Tukey's HSD intervals:

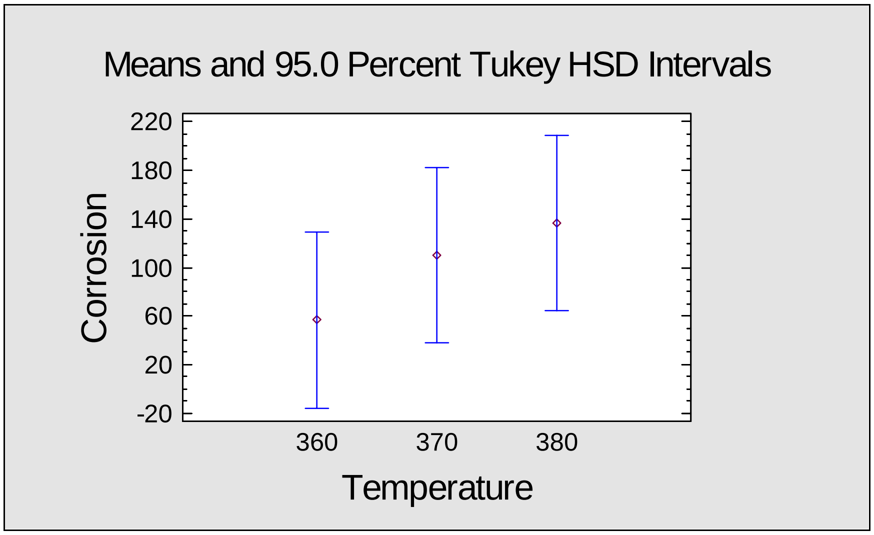

Figure 11: Means Plot for Temperature

Since all of the intervals overlap, we can not declare any means to be significantly different from any other means, which matches the insignificant result for the *Temperature* main effects in the ANOVA table. Using *Pane Options* to switch factors displays the following:

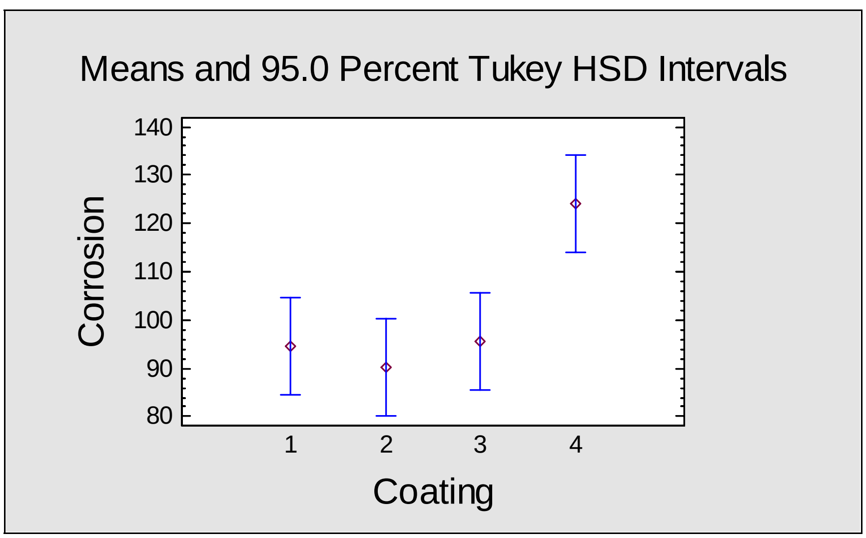

Figure 12: Means Plot for Coating

In this plot, the intervals for the first three coatings all overlap, indicating no significant differences amongst them. However, the uncertainty interval for *Coating* 4 does not overlap the other intervals, indicating a significant difference between *Coating* 4 and the other 3 coatings.

We can also select *Interaction Plot* from the *Graphs* dialog box to display the interaction between *Temperature* and *Coating*:

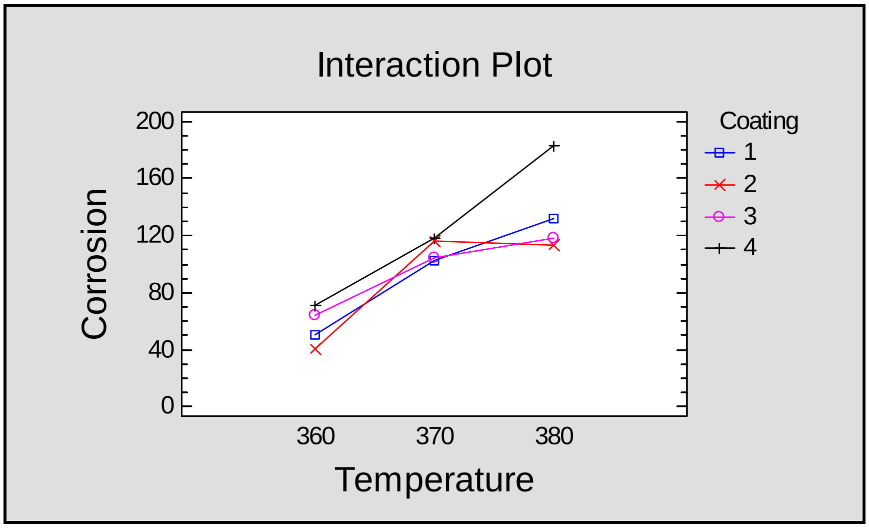

Figure 13: Interaction Plot for Temperature×Coating

Each point on the interaction plot shows the mean value for a specific combination of *Temperature* and *Coating*. Note that the difference between *Coating* 4 and the others is more pronounced at 380º than it is at the other temperatures. Also, *Coating* 2, which gives lowest *Corrosion* at 360º and 380º, is not the lowest at 370º.

We can also use *Pane Options* to add uncertainty intervals to the interaction plot:

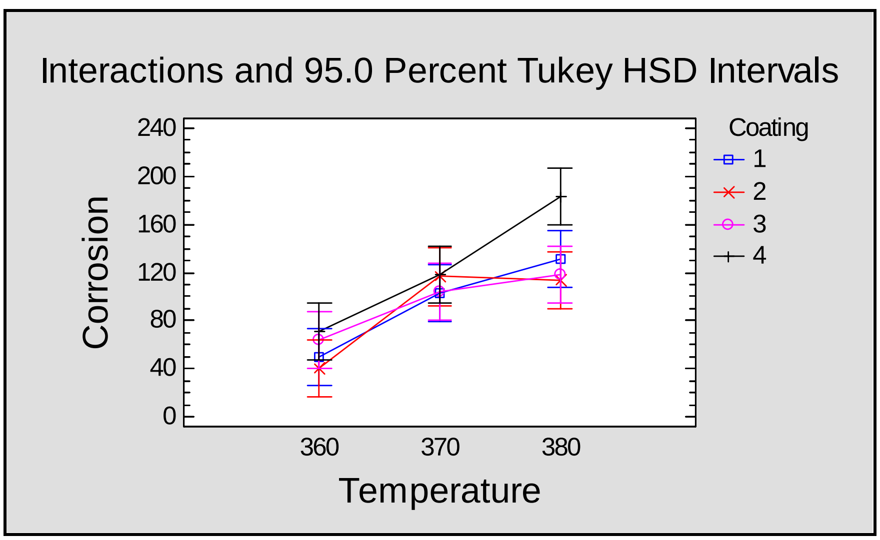

Figure 14: Interaction Plot with Uncertainty Intervals

It appears that the only temperature at which the coatings are significantly different is 380º, since all of the intervals overlap at the other temperatures.

Example #2

The second example comes from Design and Analysis of Experiments, sixth edition by Douglas Montgomery (Wiley, 2005). It is an example of an experiment performed with a single wafer plasma etching process. The goal of the experiment was to determine how *Uniformity* changed as a function of 5 factors:

A = electrode gap
B = gas flow
C = pressure
D = time
E = RF power

The first 3 factors were relatively hard to change, while the last two could be changed easily.

An excellent design for studying the effects of 5 factors is the *2^{5-1} fractional factorial design*, consisting of 16 runs at different combinations of two levels of each factor. The design is resolution V, meaning that it can estimate clearly all main effects and two-factor interactions. In running such a design, the 16 runs are ordinarily done in random order. Unfortunately, this means that all of the factors will be changed frequently, which could be expensive and time-consuming.

Montgomery describes an approach that reduces the number of times the hard-to-change factors need to be changed. He suggests arranging the 16 runs in the 2^{5-1} design as a split-plot design. In his approach, each of the 8 combinations of the 3 hard-to-change factors defines a whole-plot. Each time a whole-plot is created, two different combinations of factors D and E are tested. The layout of the design is shown below:

Whole-Plots	A	B	C	D	E	Uniformity
1	-	-	-	-	+	40.85
	-	-	-	+	-	41.07
2	+	-	-	-	-	35.67
	+	-	-	+	+	51.15
3	-	+	-	-	-	41.80
	-	+	-	+	+	37.01
4	+	+	-	-	+	91.09
	+	+	-	+	-	48.67
5	-	-	+	-	-	40.32
	-	-	+	+	+	43.34
6	+	-	+	-	+	62.46
	+	-	+	+	-	38.08
7	-	+	+	-	+	31.99
	-	+	+	+	-	41.03
8	+	+	+	-	-	70.31
	+	+	+	+	+	81.03

Figure 15: 2^{5-1} Split-Plot Experiment

Each of the 16 rows of the table represents an experiment that was performed. The – and + signs indicate whether a factor was run at its low or high level during a particular run. For example, during the first run, factors A, B, C and D were run at their low levels, while factor E was run at its high level. The measured *Uniformity* for that run was 40.85. The eight whole-plots were run in random order. Once the whole-plot conditions were set, the two runs in the whole-plot were performed, also in random order. With this approach, the hard-to-change factors A, B and C were changed less frequently then the easy-to-change factors D and E.

Step 1: Construct the Design

To construct the above experiment using STATGRAPHICS Centurion, the *Design of Experiments* section can be used. From the main menu, begin by selecting:

- If using the Classic menu: *DOE – Design Creation – Create New Design.*
- If using the Six Sigma menu: *Improve – Experimental Design Creation – Create New Design.*

On the first dialog box, select the *Screening* option (this includes the two-level factorials) and indicate that 5 experimental factors are to be varied:

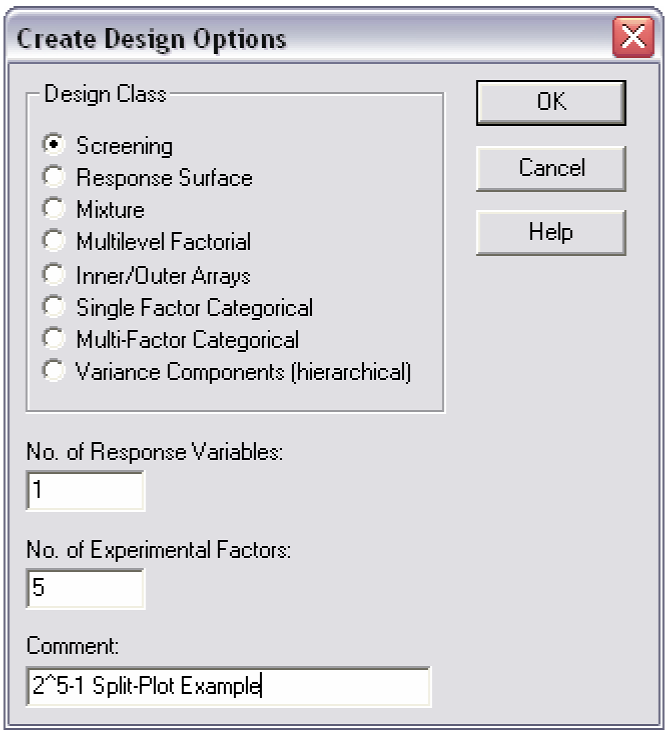

Figure 16: First Design Creation Dialog Box

On the second dialog box, specify the names of the 5 experimental factors:

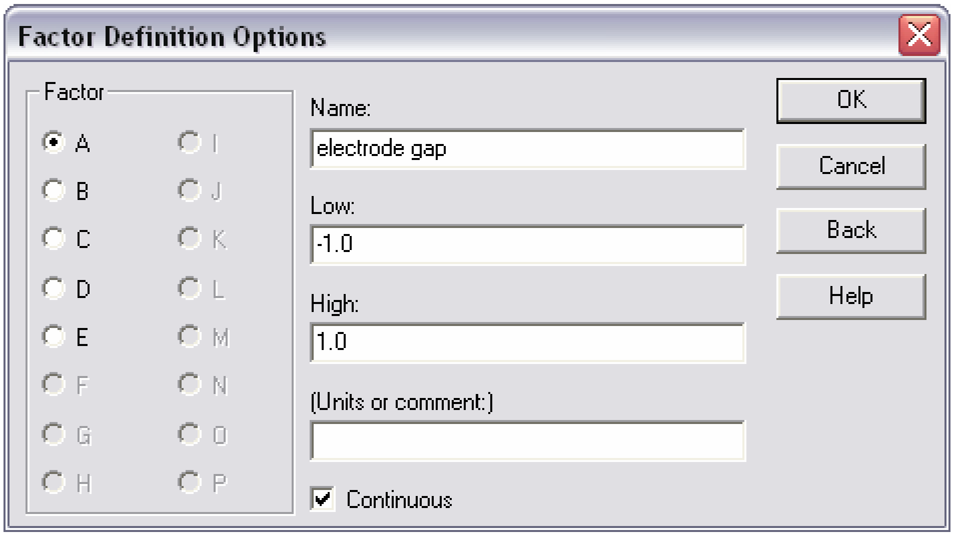

Figure 17: Second Design Creation Dialog Box

Select each button A through E and enter the following names:

A: electrode gap
B: gas flow
C: pressure
D: time
E: RF power

Since the low and high levels of each factor were not stated in Montgomery's example, leave the levels at -1.0 and 1.0.

On the third dialog box, specify a label for the response variable:

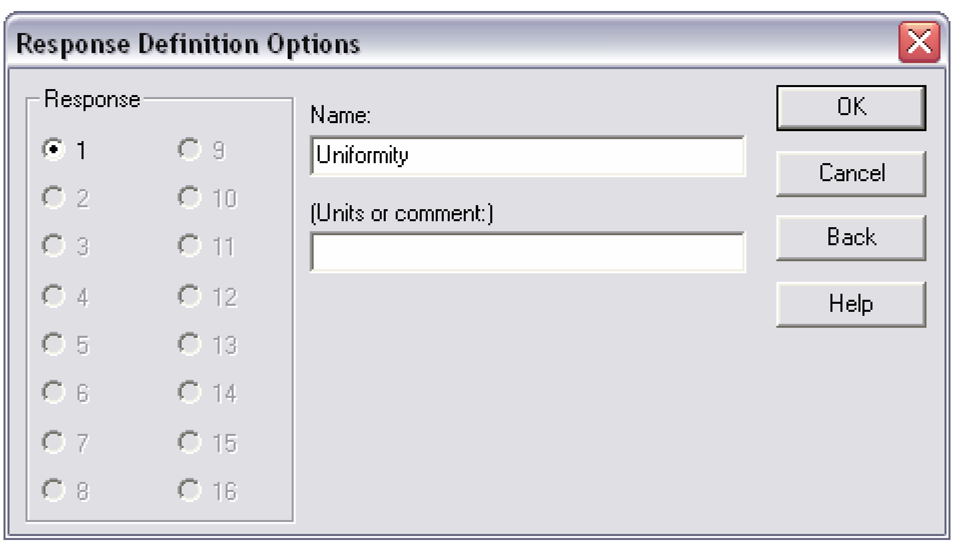

Figure 18: Third Design Creation Dialog Box

The fourth dialog box shows a list of screening designs for 5 factors:

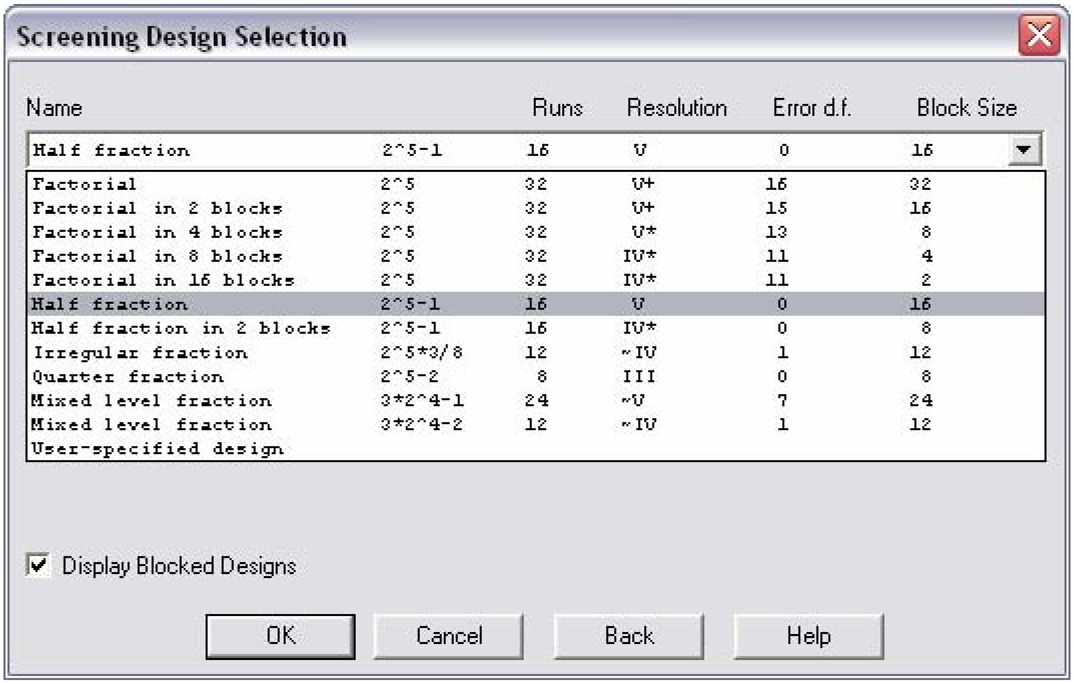

Figure 19: Fourth Design Creation Dialog Box

The design desired is the *Half fraction*. It has 16 runs and is Resolution V. Note that 0 degrees of freedom are available to estimate the experimental error. This means that all 16 runs will be used to estimate the average, the main effects, and the two-factor interactions. To obtain degrees of freedom for estimating the experimental error, additional runs could be added on the final dialog box:

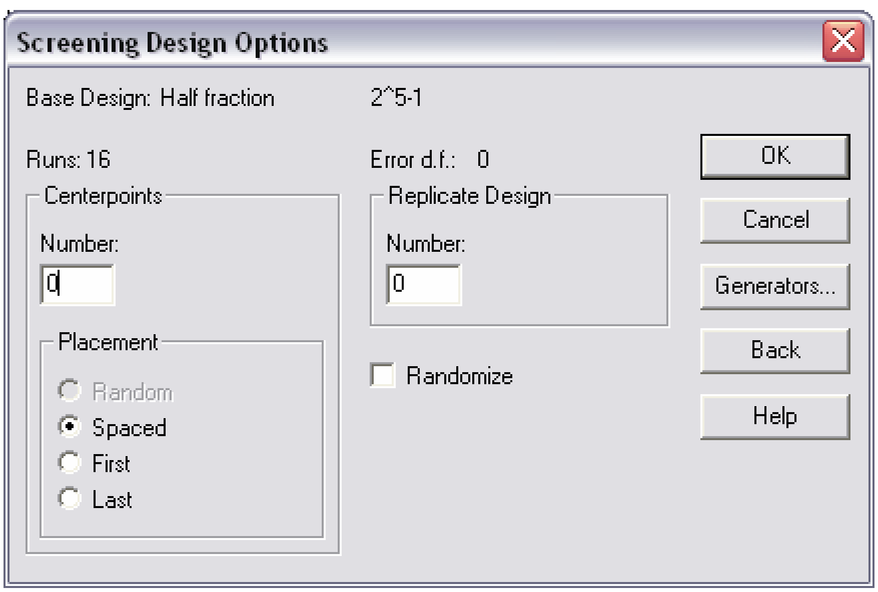

Figure 20: Fifth Design Creation Dialog Box

Most commonly, centerpoints would be added at values of each factor positioned halfway between the lows and the highs. In Montgomery's example, however, no centerpoints were used, so no formal statistical tests of significance will be possible.

After the final dialog box, the design will be created and placed into the STATGRAPHICS datasheet:

howto11a.sfx

	BLOCK	electrode gap	gas flow	pressure	time	RF power	Uniformity
1	1	-1	-1	-1	-1	1	40.85
2	1	1	-1	-1	-1	-1	35.67
3	1	-1	1	-1	-1	-1	41.8
4	1	1	1	-1	-1	1	91.09
5	1	-1	-1	1	-1	-1	40.32
6	1	1	-1	1	-1	1	62.46
7	1	-1	1	1	-1	1	31.99
8	1	1	1	1	-1	-1	70.31
9	1	-1	-1	-1	1	-1	41.07
10	1	1	-1	-1	1	1	51.15
11	1	-1	1	-1	1	1	37.01
12	1	1	1	-1	1	-1	48.67
13	1	-1	-1	1	1	1	43.34
14	1	1	-1	1	1	-1	38.08
15	1	-1	1	1	1	-1	41.03
16	1	1	1	1	1	1	81.03

A / B / C / D / E / F / G / H / I / J

Figure 21: Datasheet Containing Generated Design

The order of the design differs somewhat from that shown in Figure 15. In particular, the 8 desired whole-plots consist of the following rows:

Whole-plot #1: runs 1 and 9
Whole-plot #2: runs 2 and 10
Whole-plot #3: runs 3 and 11
Whole-plot #4: runs 4 and 12
Whole-plot #5: runs 5 and 13
Whole-plot #6: runs 6 and 14
Whole-plot #7: runs 7 and 15
Whole-plot #8: runs 8 and 16

The two runs in each whole-plot have identical values for factors A, B and C, but the levels of factors D and E are different. As stated earlier, the 8 whole-plots would be run in random order. After selecting a whole-plot, the two runs would also be done in random order.

Step 2: Analyze the Results

To analyze the results of the experiment, enter the *Uniformity* values in the rightmost column as shown in the above datasheet. Then select from the main STATGRAPHICS Centurion menu:

- If using the Classic menu: *DOE – Design Analysis – Analyze Design.*
- If using the Six Sigma menu: *Improve – Experimental Design Analysis – Analyze Design.*

Specify the response variable on the data input dialog box:

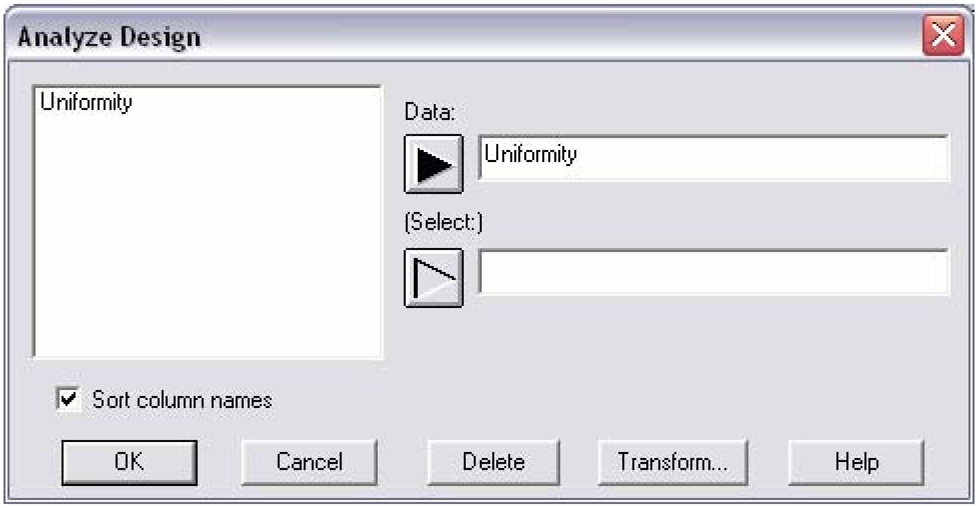

Figure 22: Analyze Design Data Input Dialog Box

An analysis window will be created, containing several tables and graphs. Of particular interest is the *Pareto Chart*, which plots the magnitude of each main effect and interaction in decreasing order:

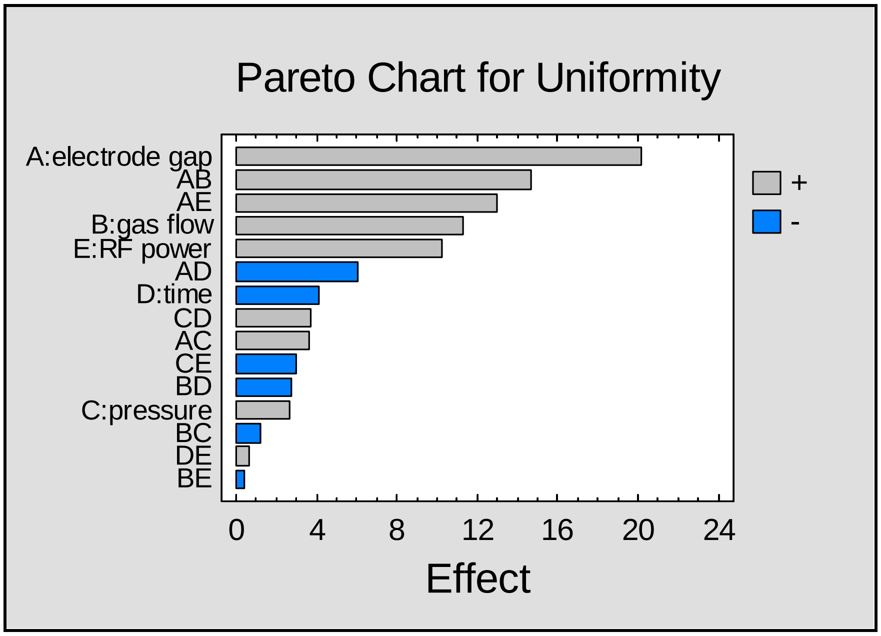

Figure 23: Pareto Chart of Estimated Effects

The largest effects appear to involve factors A, B and E and some of their interactions.

Since the whole-plot factors are subject to a different experimental error than the subplot factors, the results need to be analyzed in two groups. To analyze the whole-plot experiment:

1. Press the *Analysis Options* button  on the analysis toolbar.
2. On the *Analysis Options* dialog box, press the *Exclude* button.
3. On the *Exclude Effects Options* dialog box, double-click on any effect involving either subplot factor D or E to exclude it from the model:

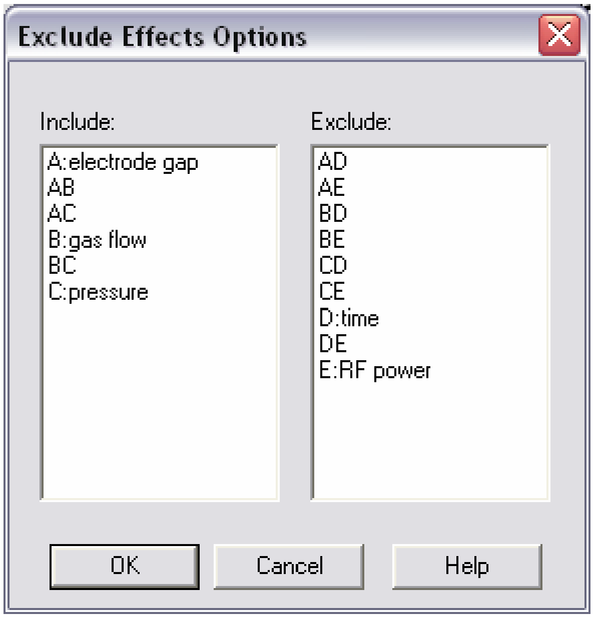

Figure 24: Excluding Subplot Effects

4. Press *OK* twice to refit the model using only factors A, B and C.
5. Press the *Graphs* button on the analysis toolbar and select *Normal Probability Plot of Effects*.
6. When the normal probability plot appears, double-click on it to maximize its pane and then press the *Pane Options* button on the analysis toolbar. Complete the options dialog box as shown below to request a *Half-Normal Plot* and to *Label Effects*:

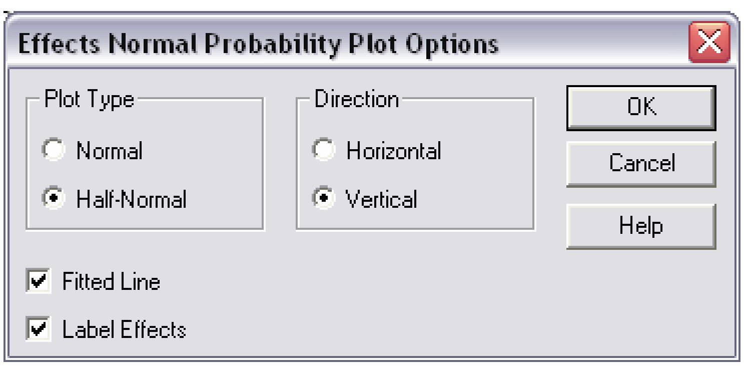

Figure 25: Normal Probability Plot Options

The above steps will create a normal probability plot that can be used to help determine which effects are statistically significant:

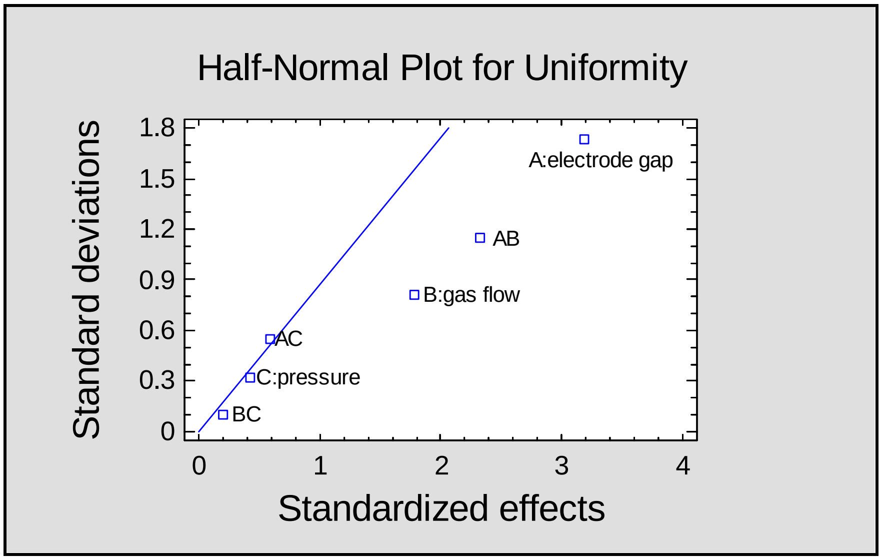

Figure 26: Half-Normal Plot for the Whole-Plot Experiment

The estimated effects that are merely manifestations of noise will appear at the bottom left and should lie approximately along a straight line. The estimated effects corresponding to real signals will lie off the line toward the upper right. In this case, it appears that both *electrode gap* and *gas flow* are statistically significant and that they interact. *Pressure* does not seem to have a significant effect.

> NOTE: the normal probability plot is the only means of determining statistical significance for the whole-plot factors in this experiment. Although the ANOVA table shows P-Values for each effect, they are based on the subplot error which does not apply to the whole-plot factors. Had degrees of freedom been available to estimate the whole-plot error, perhaps from the inclusion of several runs with replicated centerpoints, then the average result in each whole-plot could have been calculated and an analysis of variance performed on the whole-plot averages.

It is also useful in this case to display an interaction plot for factors A and B. This may be done by:

1. Press the *Graphs* button  on the analysis toolbar and select *Interaction Plots*.
2. When the plot appears, double-click on it to maximize its pane.
3. Press the *Pane Options* button and select only *electrode gap* and *gas flow*.

The resulting plot is shown below:

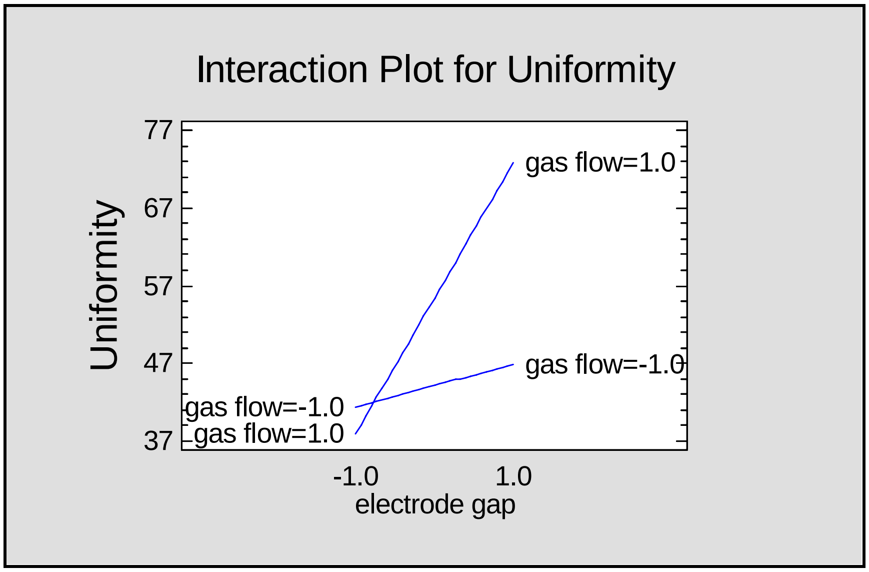

Figure 27: Interaction Plot for Electrode Gap and Gas Flow

At the low level of *electrode gap*, *gas flow* makes little difference. However, at the high level of *electrode gap*, *gas flow* makes a big difference.

It is also useful to plot the estimated model for uniformity as a function of the two important factors. The plot below was created by selecting *Response Plots* from the *Graphs* dialog box. It plots the estimated corrosion as a function of *electrode gap* and *gas flow*, with the other factors held halfway between their low and high values:

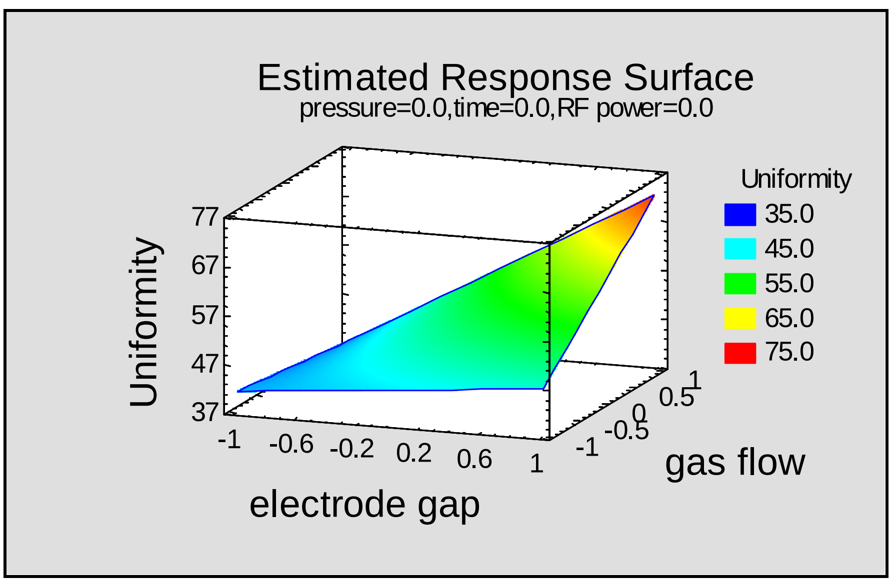

Figure 28: Uniformity versus Electrode Gap and Gas Flow

The strong interaction between the factors causes the surface to twist dramatically, resulting in unusually high *Uniformity* when both electrode gap and gas flow are at their high levels.

The same analysis can now be repeated on the subplot experiment. Returning to the *Exclude Effects Options* dialog box, we now create a model with any term (main effect or interaction) that involves a subplot factor:

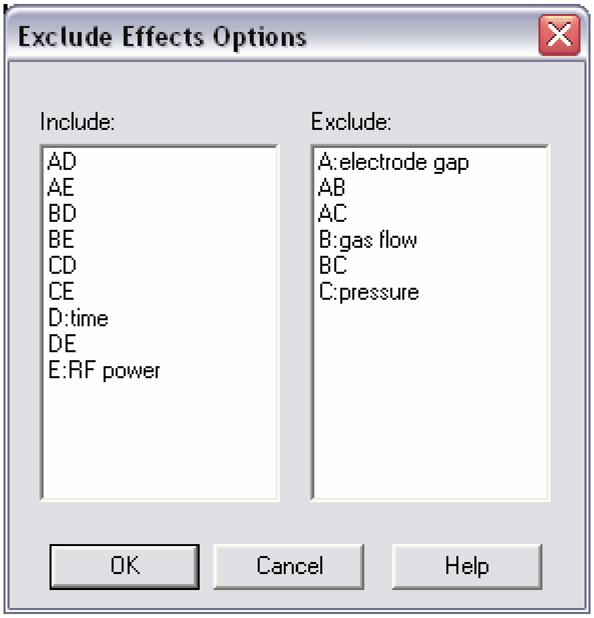

Figure 29: Excluding Whole-Plot Effects

The resulting normal probability plot shows that factor E (*RF power*) is significant, as is its interaction with factor A (*electrode gap*):

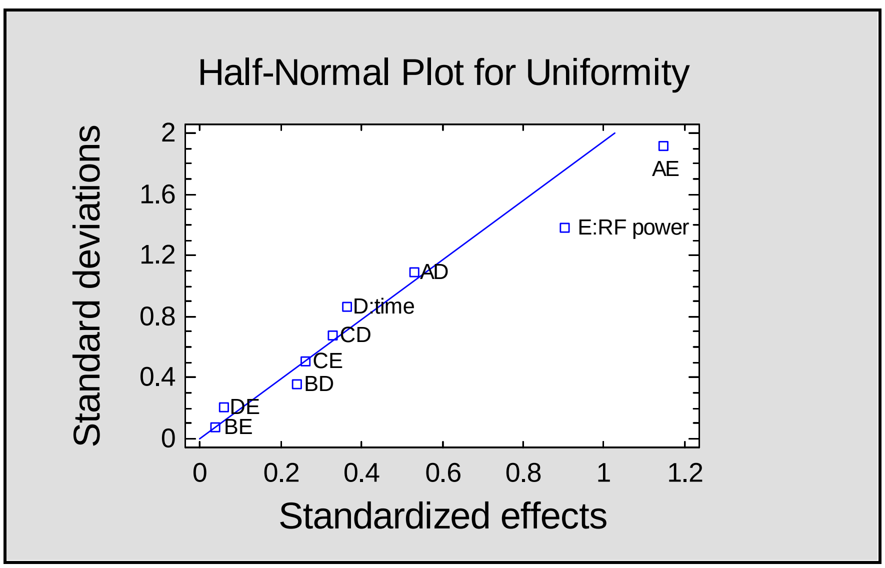

Figure 30: Half-Normal Plot for Subplot Effects

Again, an interaction plot displays the essential information:

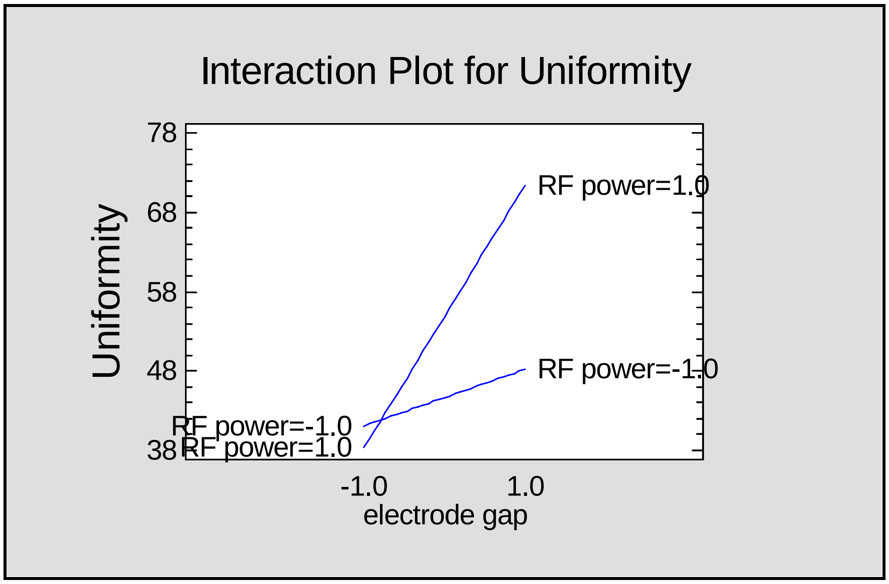

Figure 31: Interaction Plot for Electrode Gap and RF Power

As with *gas flow*, *RF Power* appears to have an effect only at the high level of *electrode gap*.

Conclusion

When the order of experiments cannot be fully randomized, a split-plot design is often useful. Those factors that cannot be changed as easily are varied across large experimental units called whole-plots, while the easily changed factors are varied across the subplots. Usually, the experimental error amongst subplots is considerably less than that amongst whole-plots, so that subplot factors and their interactions can be estimated with greater precision than the main effects of the whole-plot factors.

As pointed out by the authors of the two books referenced in this guide, it is easy to slip into a split-plot setup without realizing it. Often, the "logical" way to conduct an experiment prevents one or more factors from being fully randomized. When randomization is restricted on one or more factors, the usual tests of significance generated assuming full randomization may make those factors appear more significant than they really are. For experimenters that face the problem of not being able to fully randomize the order of their experiments, careful study of split-plot designs is a must.

Note: The author welcomes comments about this guide. Please address your responses to neil@statgraphics.com.