

Accelerated Tests as an Effective Means of Quality Improvement

Roberto G. Gutierrez, SAS Institute Inc.

ABSTRACT

Accelerated testing is an effective tool for predicting when systems fail, where the system can be as simple as an engine gasket or as complex as a magnetic resonance imaging (MRI) scanner. In particular, you might conduct an experiment to determine how factors such as temperature and voltage impose enough stress on the system to cause failure. Because system components usually meet nominal quality standards, it can take a long time to obtain failure data under normal-use conditions.

An effective strategy is to accelerate the experiment by testing under abnormally stressful conditions, such as higher temperatures. Following that approach, you obtain data more quickly, and you can then analyze the data by using the RELIABILITY procedure in SAS/QC[®]. The analysis is a three-step process: you establish a probability model, explore the relationship between stress and failure, and then extrapolate to normal-use conditions. Graphs are a key component of all three stages: you choose a model by comparing residual plots from candidate models, use graphs to examine the stress-failure relationship, and then use an appropriately scaled graph to extrapolate along a straight line. This paper guides you through the process, and it highlights features added to the RELIABILITY procedure in SAS/QC 13.1.

INTRODUCTION

Almost all reliability analysis involves answering one of two questions: “How long must you wait until a system fails?” or “How many times will a system fail within a given time period?” From a purely probabilistic standpoint, these questions are merely two sides of the same coin. How often something fails is directly related to how long you wait for each successive failure. The difference arises, however, when it comes time to model how experimental conditions (both controlled and observational) affect the outcome. To build an effective model, you must decide which question you really want answered. You answer the first question with an accelerated failure time (AFT) model that uses data from an accelerated test; the second, with a mean cumulative function (MCF) model that uses recurrent-events data. This paper is about AFT models, and it guides you through building AFT models that predict time to failure.

As an example, consider a circuit board that can operate over many possible voltages. Although higher voltages are more power-efficient, they also impose more stress on the circuit board. With higher stress comes a shorter lifespan. You might be faced with deciding whether to run the board at 1.5 volts or at 3 volts. You run an experiment in which you test several boards at each voltage.

Now, suppose that when you run your experiment, you find that your manufacturer produced high-quality boards that take a long time to fail under these “normal-use” conditions. In fact, you observed no failures at all, for the simple reason that you just ran out of time. You have no data. Although it is not necessary to observe all your boards fail, you would have to see at least some of them fail at each of the two candidate voltages.

At this point you have two choices: you can wait longer (which is often infeasible), or you can accelerate the test by increasing the voltage to 9 or even 12 volts, knowing that you are guaranteed to see some timely failures at such an unusually high voltage. Whether you just wait longer or accelerate the test, you obtain the data you need, but there is a hidden cost if you choose the higher-voltage route. If you are able to wait long enough for failures at 1.5 and 3 volts, you can simply base your decision on what you observe, with no added assumptions. You can directly estimate the median failure times for 1.5 volts versus 3 volts and decide whether the power efficiency of 3 volts is worth the (presumably) shorter lifetime. You can even get confidence limits for both median lifetimes and, more importantly, for the difference between the two.

If you instead accelerate the test by testing at 9 and 12 volts, you don't have to wait as long for your data. This can be a huge advantage, but the cost is that you no longer have the luxury of performing the direct analysis just described. Because the higher-voltage results are practically useless on their own, you must

make an assumption about the distribution of time to failure, and an assumption about how voltage affects failure time. These assumptions are what make up an AFT model.

First, you need to choose the distribution of time to failure, and you can choose from several possible distributions: exponential, Weibull, lognormal, gamma, and so on. Either you can base your choice on your substantive knowledge of the physical characteristics of your system, or you can try several distributions and choose the one that “fits” your data best. A sensible choice for circuit boards might be the Weibull distribution, because the Weibull works well for systems that consist of several working parts.

Second, you need to decide how experimental factors (voltage, in this example) affect your predicted time to failure based on the distribution that you chose previously. This amounts to choosing a *relationship function*, a function that relates the factor to the parameters of the distribution. In this example, you need to relate voltage to the parameters of the Weibull distribution that you selected in the previous step. Usually the relationship is determined by the type of stress, whether it is voltage, temperature, humidity, or another type. The relationship is based on the physical properties of stress on system components.

Analyzing data from an accelerated test is a three-step process, and the first two steps are formulating the assumptions of the AFT model as described previously. The third step is to fit the final model and extrapolate the results back to normal-use conditions. The goal is to make inference about the median (or some other percentile of) failure time for a given level of stress. How you extrapolate your results depends on which percentile, distribution, and relationship function you choose. Because results can be highly nonlinear and model-specific, graphs that transform the axes are indispensable; they make AFT analysis analogous to that for standard linear models.

Accelerated tests and their corresponding AFT models are known in the reliability literature under the name *accelerated life testing* (ALT) or the equivalent *highly accelerated life testing* (HALT). Both Nelson (1990) and Meeker and Escobar (1998) provide complete treatments of ALT methods, and Escobar and Meeker (2006) give a more concise review.

The next section describes how to interpret AFT models. The subsequent sections cover the three steps in AFT analysis.

THE BASICS OF AFT MODELS

Accelerated failure time (AFT) models are a broad class of parametric regression models that depict how one or more factors affect the reliability of a system. These factors serve to either accelerate or decelerate the accumulation of risk that ultimately results in failure.

For a system with only one factor that affects failure, consider the reliability (or survivor) function $S_0(t)$, which represents baseline reliability at a given time t , where “baseline” means that the factor is kept at some nominal or reference level (usually 0). This function is the probability that the reference system does not fail before time t ; it is equal to 1 at time $t = 0$ and decreases to 0 as time elapses. The rate at which reliability decreases with time is not constant and depends on the properties of the failure time distribution.

In an AFT model, failure is accelerated by multiplying time by a constant factor. That is, the reliability of a system with a factor value of x (not the reference value) is $S_x(t) = S_0(\theta t)$, where θ is a positive constant that depends on x . If $\theta < 1$, time passes more slowly for this system (compared to the reference system), and failure will take longer. Conversely, if $\theta > 1$, time passes more quickly for this system, and failure will occur sooner. Whereas functional forms of $S_0(t)$ and its related $S_x(t)$ are derived from the failure time distribution, the mapping from x to θ derives from the relationship between stress and failure as assumed by the AFT model.

This interpretation of an AFT model applies to any choice of $S_x(t)$. However, the real question—what is the useful life of my system when the factor is equal to x ?—is muddled by the mathematics. To make things clearer, $S_x(t)$ is often chosen to be from a *location-scale family*, such as the lognormal or Weibull family. For a distribution from a location-scale family, from $S_x(t)$ you can derive a simple form for a given percentile of failure time, which is what you are interested in anyway. Formally,

$$\log\{t_p(x)\} = \mu(x) + \sigma Q(p)$$

where $t_p(x)$ is the $(100 \times p)$ percentile, $\mu(x)$ is the location, σ is the scale, and $Q(p)$ is the $(100 \times p)$ percentile of some standard distribution, such as the normal.

The percentile of interest, t_p , depends on x through the location function $\mu(x)$, and $\mu(x)$ often takes the form

$$\mu(x) = \beta_0 + \beta_1 g(x)$$

for some relationship function $g(x)$. That is, you choose a relationship function that transforms the log-percentile's dependence on x into a linear function. For simplicity, you usually hold the scale σ constant, but you could also make that depend on x . That your model is expressed in terms of log-percentiles and not just percentiles shouldn't bother you—it is easy to transform between the two.

INSULATION DATA

Hahn, Meeker, and Doganaksoy (2003) analyze data from an accelerated test of a new insulation design for generator armature bars (GABs). The insulation is known to degrade at a faster rate the higher the voltage. The engineers were interested in the 1st and 5th percentiles of insulation lifetime at 120 V/mm. "V/mm" refers to voltage adjusted for thickness across a dielectric, which we refer to simply as voltage in what follows.

The following SAS statements create the data set GAB, which contains data about endurance tests of 75 electrodes, 15 at each of 5 voltages that are greater than 120 V/mm:

```
data GAB;
  input Voltage Hours Censor Count @@;
  datalines;
170 6.480 1 15      190 3.248 0 1      190 4.052 0 1      190 5.304 0 1
190 6.480 1 12      200 1.759 0 1      200 3.645 0 1      200 3.706 0 1
200 3.726 0 1       200 3.990 0 1      200 5.153 0 1      200 6.368 0 1
200 6.480 1 8       210 1.401 0 1      210 2.829 0 1      210 2.941 0 1
210 2.991 0 1       210 3.311 0 1      210 3.364 0 1      210 3.474 0 1
210 4.902 0 1       210 5.639 0 1      210 6.021 0 1      210 6.456 0 1
210 6.480 1 4       220 0.401 0 1      220 1.297 0 1      220 1.342 0 1
220 1.999 0 1       220 2.075 0 1      220 2.196 0 1      220 2.885 0 1
220 3.019 0 1       220 3.550 0 1      220 3.566 0 1      220 3.610 0 1
220 3.659 0 1       220 3.687 0 1      220 4.152 0 1      220 5.572 0 1
;
```

The variable **Voltage** is the voltage, the variable **Hours** is the insulation lifetime in thousands of hours, and the variable **Censor** equals 1 if the insulation did not fail by the end the experiment and equals 0 if **Hours** is the failure time. The variable **Count** is a frequency variable that represents duplicated observations. Note that at a voltage of 170 V/mm, no failures were observed and all 15 observations are censored at 6,480 hours.

STEP 1: ESTABLISH A DISTRIBUTION

The first step in an AFT analysis is establishing a failure time distribution. PROC RELIABILITY offers several choices, including the exponential, log-logistic, lognormal, and Weibull distributions. You can base your choice either on the physical properties of the system or on your experience with similar systems. For example, an exponential distribution assumes constant risk and therefore is seldom appropriate in engineering. The Weibull and lognormal distributions are commonly used in engineering, and are appropriate for systems where there are many working parts or, as is the case in this example, where failure can occur in any of numerous areas within the system. The log-logistic distribution is very similar to the lognormal, and the two are often used interchangeably. For more information about these and other available distributions, see the chapter about the RELIABILITY procedure in the *SAS/QC 13.1 User's Guide* (SAS Institute Inc. 2013).

In the absence of any physical knowledge or experience, you can try several distributions and choose the one that has the best empirical fit. You can examine the fit by using a *probability plot*. A probability plot is a plot of the data along with the fitted distribution, for each stress level, with the axes scaled in a way that makes evaluating the model fit easy.

The following SAS statements produce location and scale estimates for each voltage level, tables of estimated percentiles (and standard errors and confidence limits), and a probability plot.

```

proc reliability data = GAB;
  distribution weibull;
  freq Count;
  probplot Hours*Censor(1) = Voltage / overlay noconf;
run;

```

The DISTRIBUTION statement sets the distribution as Weibull, and the FREQ statement establishes **Count** as the frequency variable. The PROBLOT statement performs the estimation and generates the probability plot.

Because you have not yet fit a regression model, estimation is performed separately at each voltage level. [Figure 1](#) provides the Weibull parameter estimates for 190 V/mm. PROC RELIABILITY also produces tables for the other voltages, although they are not presented here. At 190 V/mm, the location parameter is 2.49 and the scale parameter is 0.42.

Figure 1 Weibull Parameter Estimates for 190 V/mm

The RELIABILITY Procedure					
Weibull Parameter Estimates					
Parameter	Estimate	Standard Error	Asymptotic Normal		Group
			95% Confidence Limits		
			Lower	Upper	
EV Location	2.4923	0.4385	1.6328	3.3518	190
EV Scale	0.4220	0.2349	0.1417	1.2566	190
Weibull Scale	12.0887	5.3013	5.1180	28.5535	190
Weibull Shape	2.3695	1.3190	0.7958	7.0547	190

The location and scale are in terms of the extreme-value (EV) distribution, which is the distribution of log-time when time follows a Weibull distribution.

[Figure 2](#) provides a table of estimated percentiles, standard errors, and confidence limits for 190 V/mm.

Figure 2 Weibull Percentiles for 190 V/mm

The RELIABILITY Procedure					
Weibull Percentile Estimates					
Percent	Estimate	Standard Error	Asymptotic Normal		Group
			95% Confidence Limits		
			Lower	Upper	
0.1	0.65518004	0.83961262	0.05315531	8.07559718	190
0.2	0.87800715	0.98507151	0.09738696	7.91580844	190
0.5	1.29335866	1.18040953	0.21619863	7.73722117	190
1	1.73470533	1.31218059	0.39386971	7.64009637	190
2	2.32916043	1.40555013	0.71373198	7.60087607	190
5	3.45125369	1.42458525	1.53682227	7.75050717	190
10	4.67641213	1.37371239	2.62947686	8.31679896	190
20	6.41884933	1.56600317	3.9791472	10.3543862	190
30	7.82385968	2.13343854	4.58471472	13.3514916	190
40	9.10455597	2.90972887	4.86656928	17.0331366	190
50	10.3562107	3.83230952	5.01434558	21.388853	190
60	11.65078	4.91211682	5.09891322	26.6214915	190
70	13.0737902	6.21541546	5.14910171	33.1949145	190
80	14.7775485	7.90755802	5.17748869	42.1779656	190
90	17.1888396	10.5061455	5.18771368	56.9530678	190
95	19.2079037	12.8380296	5.18275277	71.1867964	190
99	23.0299074	17.5793584	5.15877714	102.810534	190
99.9	27.3280801	23.3420547	5.12344795	145.765892	190

The median of the minimum extreme-value distribution is approximately -0.37 , which corresponds to an estimated log-median time of

$$\log(t_{0.50}) = 2.49 + 0.42(-0.37) = 2.33$$

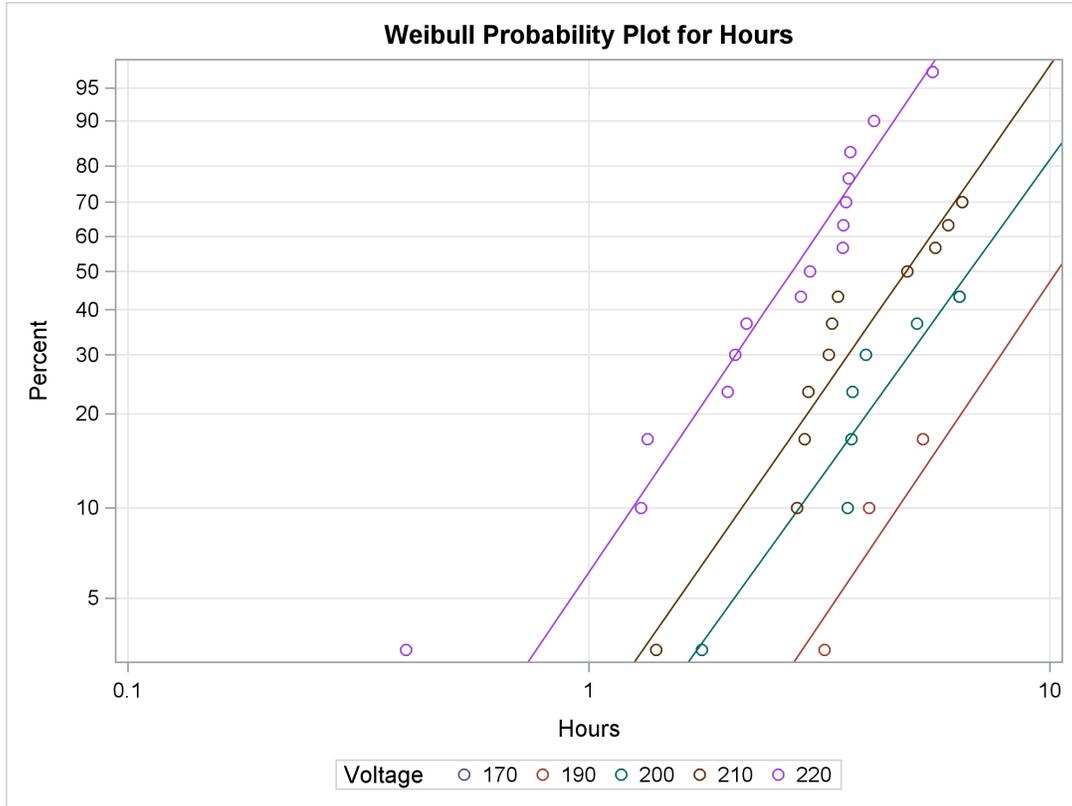
The estimated median failure time is then

$$t_{0.50} = \exp(2.33) = 10.28 \text{ thousand hours}$$

Within rounding error, this corresponds to the estimate in [Figure 2](#).

[Figure 3](#) shows the probability plot and overlays the Weibull fits at all voltage levels (except 170 V/mm, where all observations are censored).

Figure 3 Weibull Probability Plot



The **Hours** axis is log-transformed and the probability axis is Weibull-transformed, making the probability curves linear. Except for one possible outlier, the Weibull model provides a very good fit to the data. As an exercise, try the lognormal distribution, which you will find to be an equally good fit.

STEP 2: ESTABLISH A RELATIONSHIP FUNCTION

After you have established a Weibull distribution for these data, the second step is to fit a model that relates voltage stress to failure time. In the previous analysis, you made no assumption about how stress affects failure; instead you fit separate Weibull models for each voltage level.

Stress-life relationships are almost always determined by the physical properties of the system, and the two most commonly used are the Arrhenius and the inverse-power relations. The Arrhenius relation is a reciprocal function appropriate for stress from either temperature or activation energy. The inverse-power relation is a log transformation appropriate for a wide variety of physical stresses: voltage, pressure, vibration, weight load, and so on. For more information about the physical principles that govern these relations, see Nelson (1990).

The following SAS statements produce regression parameter estimates and a diagnostic plot for the Weibull/inverse-power AFT model:

```
proc reliability data = GAB;
  distribution weibull;
  freq Count;
  model Hours*Censor(1) = Voltage / relation = power;
  relationplot Hours*Censor(1) = Voltage /
    fit = model
    plotdata
    plotfit;
run;
```

The MODEL statement fits a regression model that has an inverse-power stress relationship, called “power” in PROC RELIABILITY. The RELATIONPLOT statement graphs the relationship between stress and median lifetime (median is the default) and overlays the data onto the model fit.

Because the inverse-power relationship is a log transformation, Figure 4 establishes the linear relationship between log-percentile and voltage as

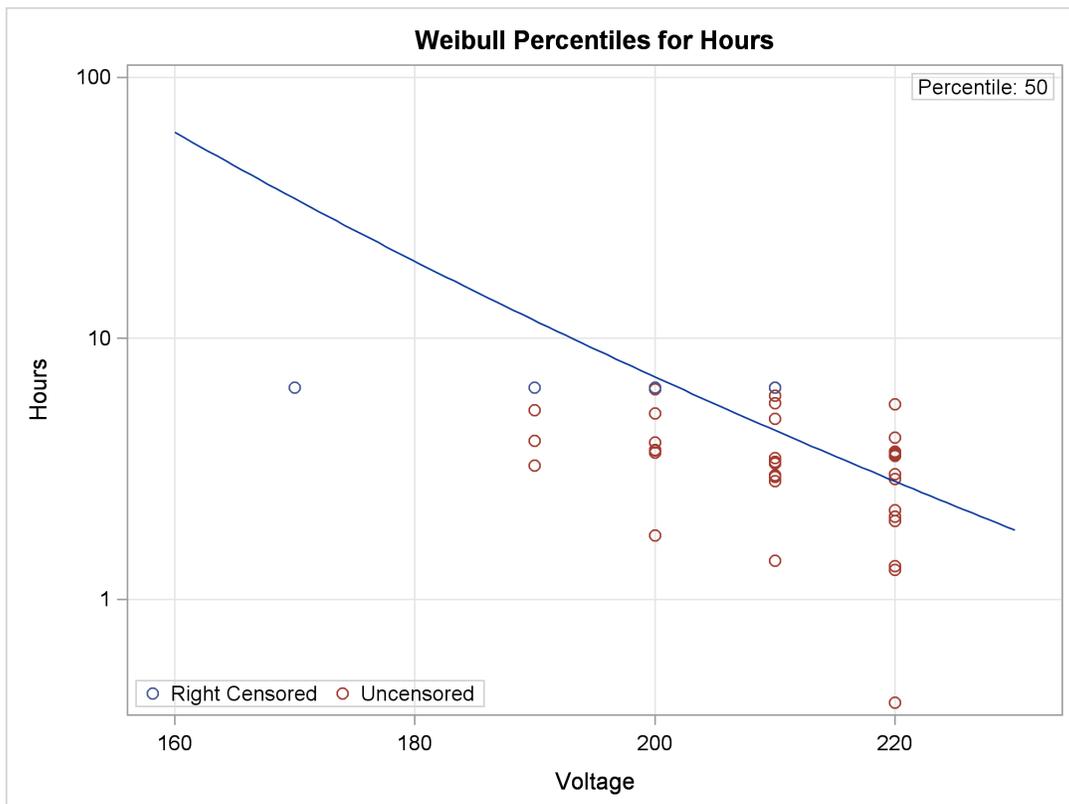
$$\begin{aligned} \log(t_p) &= \beta_0 + \beta_1 \log(\text{voltage}) + \sigma Q(p) \\ &= 53.39 - 9.68 \log(\text{voltage}) + 0.44 Q(p) \end{aligned}$$

Figure 4 Regression Parameters for the Weibull/Inverse-Power Model

The RELIABILITY Procedure				
Weibull Parameter Estimates				
Parameter	Estimate	Standard Error	Asymptotic Normal	
			95% Confidence Limits	
			Lower	Upper
Intercept	53.3871	9.4614	34.8431	71.9310
Voltage	-9.6753	1.7683	-13.1412	-6.2095
EV Scale	0.4402	0.0616	0.3346	0.5791
Weibull Shape	2.2717	0.3178	1.7270	2.9883

Figure 5 is a plot of voltage versus median time ($Q(.5) = -0.37$), with the observed data overlaid. At lower voltages, many observations are censored, so the median time is not observed within the data. For this reason, diagnosing model fit can be tricky at lower voltages. At the two highest voltages, however, the fit is quite reasonable.

Figure 5 Relation Plot, Weibull/Inverse-Power Model



STEP 3: EXTRAPOLATE TO NORMAL-USE CONDITIONS

When you are satisfied with the model fit, you can extrapolate the results of the accelerated test back to normal-use conditions. Recall that in this example, the engineers were interested in the 1st and 5th percentiles of failure time at 120 V/mm. To estimate these percentiles at this stress level, modify the code from step 2 as follows:

```
proc reliability data = GAB;
  distribution weibull;
  freq Count;
  model Hours*Censor(1) = Voltage / relation = power;
  relationplot Hours*Censor(1) = Voltage /
    fit = model
    plotdata
    plotfit 1 5
    relation = power
    slower = 100
    supper = 240
    sref(intersect) = 120
    sreflabel = "Normal Use";
run;
```

The "PLOTFIT 1 5" option specifies plots of the desired percentiles instead of the default median. The RELATION option specifies a stress-axis transform that corresponds to the one used in the regression model; this makes the percentile curves linear. The SLOWER and SUPPER options control the range of the stress axis, ensuring that the normal-use voltage is included. The SREF and SREFLABEL options draw reference lines that make it easy to read the estimated percentiles of interest. These enhanced reference lines are a new feature of SAS/QC 13.1.

Figure 6 Extrapolation Plot, Weibull/Inverse-Power Model

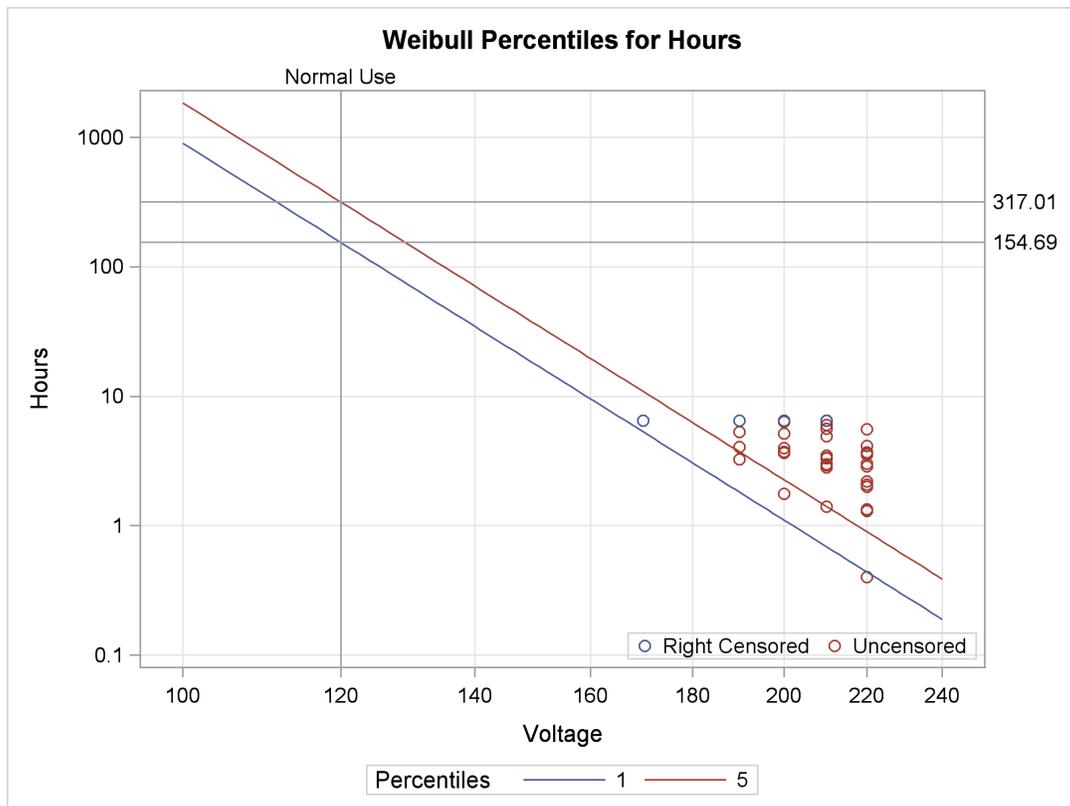


Figure 6 shows the resulting plot. In normal-use conditions, the estimated 1st percentile is 154,690 hours and the estimated 5th percentile is 317,010 hours. The insulation lasts a very long time under normal stress.

SUMMARY

Accelerated tests are often used when establishing lifetime under normal-stress conditions proves infeasible because of time constraints. Although it is convenient, relating test results to normal use requires specifying both a distribution of lifetimes and a stress-lifetime relationship. Knowing the physical characteristics of the system, empirical evidence, and trial and error can all contribute to building an effective model. Model diagnostics and extrapolation to normal-use conditions are made easier by graphs, as provided by PROC RELIABILITY in SAS/QC 13.1.

REFERENCES

- Escobar, L. A. and Meeker, W. Q. (2006), "A Review of Accelerated Test Models," *Statistical Science*, 21, 552–577.
- Hahn, G. J., Meeker, W. Q., and Doganaksoy, N. (2003), "Accelerated Testing for Speedier Reliability Analysis," *Quality Progress*, 36, 58–63.
- Meeker, W. Q. and Escobar, L. A. (1998), *Statistical Methods for Reliability Data*, New York: John Wiley & Sons.
- Nelson, W. (1990), *Accelerated Testing: Statistical Models, Test Plans, and Data Analyses*, New York: John Wiley & Sons.
- SAS Institute Inc. (2013), *SAS/QC 13.1 User's Guide*, Cary, NC: SAS Institute Inc.

ACKNOWLEDGMENTS

The author is grateful to Bob Rodriguez of the Advanced Analytics Division at SAS Institute Inc. for his valuable assistance in the preparation of this manuscript.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author:

Roberto G. Gutierrez
SAS Institute Inc.
SAS Campus Drive
Cary, NC 27513
bobby.gutierrez@sas.com

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.