Predict Failures:
Crow-AMSAA 101 and Weibull 101

Paul Barringer, P.E.
Barringer & Associates, Inc.
P.O. Box 3985
Humble, TX 77347-3985
USA

Phone: 1-281-852-6810
FAX: 1-281-852-3749

Email: hpaul@barringer1.com
Web: http://www.barringer1.com

International Mechanical Engineering Conference
http://www.imece2004.com

Kuwait
December 5-8, 2004

Proceedings of IMEC 2004
International Mechanical Engineering Conference
December 5-8, 2004, Kuwait

IMEC2004-FM201-CP

Predict Failures: Crow-AMSAA 101 and Weibull 101

Paul Barringer, P.E.

Barringer & Associates, Inc.
P.O. Box 3985
Humble, TX 77347-3985 USA
Email: hpaul@barringer1.com

Abstract

Reliability growth plots known as Crow-AMSAA plots are powerful for predicting future failures for *mixed failure modes*. Weibull probability plots are powerful *single failure mode* tools for predicting the type of failure mode which guides reliability centered maintenance strategies and forecasting future failures for each failure mode. Both analytical tools are minimum requirements for every reliability engineer's the tool box. Real data examples are shown to illustrate the value for acquiring engineering/maintenance data and using these basic tools to give the data a voice. Giving data an unemotional voice is a rational for decisions and for corrective action. Examples and illustrations describe the basics of each tool.

Keywords: Crow-AMSAA plots, reliability growth plots, Weibull probability plots, Weibull analysis, failure forecast, failure modes, data, maintenance data, reliability analysis.

Data

Most chemical plants, refineries, and manufacturing systems have objective, variables type data in computerized maintenance management (CMMS) systems. The data is used in two ways:
1) cumulative time vs cumulative failures for Crow-AMSAA (C-A) plots and 2) age-to-failure for Weibull plots. Data acquisition is simpler for C-A plots than for Weibull plots. Weibull plots provide smarter information than C-A plots, such as failure modes; and the Weibull methodology allows forecast of future failures knowing only the age of survivors and statistics about the way the population fails.

For C-A plots, the X-axis is often simply chronological time which is a cumulative number. When the fleet size is changing (up or down) accumulate time in cumulative fleet time, (e.g., cumulative fleet hours) so as to get the cumulative time events in the correct relationship. C-A plots require cumulative time versus cumulative failures (or cumulative events). They make reliability visible.

Weibull data must be presented in the format of age to failure. You must know the birth date of the component (time zero) and you must know the death date of the component with the elapsed time resulting in an important statistic called age-to-failure. In short, you must produce data as occurs with humans, as their death certificates usually state the chronological age and reason(s) for death. For humans, we report life in years although for equipment the age-to-failure can be years, cycles, burn-out cycles, hours in service plus a multiplier times the number of burn-out cycles, hours of use plus hours of idleness, etc. for motivation toward failures.

CMMS systems must reduce the mass of data into simple statistics for rational subsets of data with ages-to-failure, for example the loads, sizes and types of bearings have considerable variability and cannot be lumped together for details analysis; although at a very high altitude lumping data together with simple arithmetic mean times to failure (MTTF) for non repairable items and mean times between failure (MTBF) for repairable items may be as helpful as a moisten finger held vertically to find the direction of the wind.

Some items do not die (fail) in service and contain censored information about their death date. Some items are removed without failure after acquiring some age and we can never know how much remaining life existed in the item as equipment fails on a probabilistic basis rather than a deterministic basis. Censored data (or the synonym for censored data is suspended data) is useful for reliability calculations, therefore do not discard the suspended data as it contains valuable information. Also failures from an event rather than an aging event results in censored data. This means you must record the censored data in your CMMS along with reasons for removal. More will be said below about suspended data for Weibull analysis. Censored data can also exist for C-A plots and the censored information will be contained only in the last data point in a C-A plot which contains time but it has not yet reached the next failure.

To acquire failure data, you must define what constitutes a failure. Without a rational definition of failure, each person will record something different, and this will result in data chaos and inconsistencies. Without a well understood definition of failure, arguments ensue which generates much heat but little light will be shed on the issues. Modern medicine has shown the value of attaching specific names to illnesses rather than just reporting the patient is sick. Consider your opinion of your doctor if your personal medical record simply reported you visited the doctor on a certain date and the cause was reported only as “sick”----you would view your personal medical history as wholly ineffective for building your medical case history. The lack of specificity for human failures (death) over the past centuries makes building a family’s genetic problems very difficult to accomplish even when death records for the family have been carefully maintained. The bottom line is this: Collecting and using data from CMMS must be carefully considered and preplanned by thoughtful analysis to prevent accumulating non-useable data garbage.

All data in a CMMS have accuracy problems particularly in the age to failure. Lack of accuracy in age to failure is a common deficiency. For example, tubes in a heat exchanger may only be diligently inspected at year 3 (1 tube found leaking), followed by year 6 (5 tubes found leaking), followed by year 9 (17 tubes found leaking), followed by year 11 (42 tubes found leaking). At precisely what age did the heat exchanger tubes fail?—we do not know because we can only discover failures at inspection intervals. Unfortunately, seldom are inspections performed on the remaining 600 tubes when the heat exchanger was retubed after 12 years of service and thus we did not harvest critical information for future decisions (likewise we lack the detailed facts to show wise decisions were made for removing the heat exchanger from service!). A huge amount of tube failures are never disclosed at the oldest age (in this case, 12 years) because an autopsy was not performed on tubes retired. After all, few tubes are found in the failed condition at the early ages; but you should expect to find many tubes failed at old age and this end of life information would be very valuable for historical records to forecast future failures. Thus with equipment as with humans, autopsies at end of life provide valuable data for future conditions. Don't miss the opportunity to acquire more data and better data by inspecting retired equipment.

For emphasis, we cannot know precisely on what date the first heat exchanger tube failed (it failed some time between time zero and three years of time). Nor can we know the individual failure dates for any of the other ages to failure. Of course, we're blind for how many failures existed at age 12 years when the heat exchanger was removed from service and retubed. Unless the heat exchanger is inspected and failures recorded at year 12, we must extrapolate---in fact the unreported failures at age 12 are very important for future events as the aging process says we should expected more deaths per year in the period between 11 and 12 years. Just as occurs with human deaths, autopsies are important events prior to burial and the same is true for death of equipment to set the record straight for future events.

We have two deficiencies in data for this heat exchanger: 1) we only know deaths occurred over an interval, and 2) we're blind as to the number of deaths within the interval from 11 years to 12 years where data on the old veteran tubes strongly influence the results for predicting future conditions. Perhaps we can never, in a practical fashion, acquire accurate ages to failure on the heat exchanger tubes, but we can perform an inspection upon removal of the system to pinpoint the deaths occurring in the short interval between year 11 and year 12 to give the most accurate time and where we should have expected more failures to occur per unit of time so as to meliorate data deficiencies.

In general, we do not record failures to accurate time scales unless life data is acquired in the laboratory where precise failure criteria have been established and accurate time keeping equipment has been installed. For field failures we may get time recorded to the hour, but more frequently we're at best accurate to the day, sometimes we're only accurate to the week or month, and as the heat exchanger example illustrates, unknowns in time may be recorded in years. Many people recognize we need more accurate data to achieve more accurate analysis but like the weather, we only talk about doing something with making improvements.

Typically, reliability data has many errors in recording ages-to-failure. Customarily, Weibull analysis recognizes that errors in the X-direction (the time scale) are greater than errors in the Y-direction (the probability scale) on Weibull plots. Consequently the Weibull data is regressed X-onto-Y recognizing the time axis (X) has the largest errors compared with the median rank plotting positions associated with the Y-axis. Customarily, C-A plots do not recognize the error condition in the time scale and consequently the customary math technique regresses Y-onto-X.

The sad part about failure data is this: Most engineers don't know how to use the data they already have for making better decisions. We seem to operate in a perpetual state of oblivion regarding data! In short: We just don't get it!

Seventy years ago we accepted that pilots of airplanes could fly in fair weather when they could see the ground. Few flights were involved and flight by the seat of our pants was acceptable. Today we reject seat of the pants flight, and we demand that commercial airline pilots must fly their planes, navigate, and follow precise rules by the numbers. Similarly today we expect medical doctors must be guided by the numbers for laboratory test and careful diagnosis of human ills to find a solution in diagnosis, surgery, and pharmacology with medicine by the numbers. Likewise we even demand that simple disinfection of water and treatment of sewage be performed by the numbers.

Most engineers need to understand they are a dying breed if they plan to maintain plants and equipment using only the seat of their pants, qualitative data, for making decisions. Engineers must fluently use reliability data so they can reduce costs and avoid failures by using data to make wise decision. Without the numbers, we engineers will soon be viewed as technical amateurs with declining pay scales and unemployment as the byproduct. Engineers—wake up! Use the data in your maintenance systems to solve technical problems and make improvements! The task of reliability engineers is to avoid failures which carry a requirement for solving problems with data. The task of maintenance engineers is to quickly restore equipment to operating conditions which requires understanding failure modes and the failure data. Both reliability engineers and maintenance engineers can use a common set of data for an excellent communication tool to solve the vitally few problems in the shortest interval of time using facts from the data to reduce costs for our manufacturing plants.

Lack of clues, on how to use the data, results in poorer financial performance for our businesses. We often incur higher cost by not giving the data a voice by use of Crow-AMSAA plots or by Weibull plots. Frequently failure patterns give clues for avoiding future failures, and we cannot see the forest for the trees unless we use our analytical tools. Does all data need analysis?—no! Some data on low cost failures does not warrant they higher cost of analysis. Does high failure costs need analysis?—yes! We need to mitigate the failures by bringing an arsenal of tools into play to reduce costs. Give the data a voice to solve the vital few economic problems driven by the facts and not emotions. Make the data talk!

Data facts begin with a Pareto distribution of economic problems measure in money (not a nose count of problems as this establishes the problem solving priority). Data and supporting details are often found in maintenance databases Convert the detailed information into an understandable model for solving the problem. Consider the data in Table 1 which can be used two ways: 1) Crow-AMSAA plots of mixed failure modes and 2) Weibull plots of individual failure modes.

Crow-AMSAA Plots 101-

C-A plots are simple power curves of cumulative time and cumulative failures. The data, when plotted on log-log paper, usually results in straight lines. Cumulative failures are plotted on the Y-axis. Cumulative time is plotted on the X-axis. The math is simple for the regression line, using the equation $N(t) = \lambda * t^{\beta}$, and provides two statistics (line slope, β, and Y-axis intercept at time t=1, λ). Trend line slope (beta) statistic is a powerful indicator of increasing, decreasing, or a state of no improvement/deterioration. The Y-intercept provides the failure rate at time equal to 1 which is simply a hypothetical value of major interest to allow forecasting of future failures.

C-A plots make reliability visible even with mixed failure modes and often without starting the data acquisition at time zero.

Table 1: Raw Data From Maintenance Database

Equipment Log Equipment runs 24 hours per day 7 days/week						
Date	**Time**	**Event**	**Maint. Time**	**Cum. Age (hrs)**	**Failure Age (hrs)**	**Comment**
1/2/88	8:00	Start-up		0.0		Commissioned
1/3/88	1:00	Right Shaft Brg Down		17.0	17.0	Bearing failure
1/4/88	7:00	Repaired & Up	30.0	17.0		
1/4/88	11:00	Right Shaft Brg Down		21.0	4.0	Bearing failure
1/4/88	18:00	Repaired & Up	7.0	21.0		
1/10/88	17:00	Right Shaft Brg Down		164.0	143.0	Bearing failure
1/11/88	2:30	Repaired & Up	9.5	164.0		
1/14/88	20:00	Coupling Down		253.5	253.5	Coupling failure
1/15/88	10:30	Repaired & Up	14.5	253.5		
11/29/88	3:00	Main Shaft Seal Down		7902.0	7902.0	Main Seal Failure
11/29/88	14:30	Repaired & Up	11.5	7902.0		
6/3/89	12:00	Main Shaft Seal Down		12363.5	4461.5	Main Seal Failure
6/4/89	3:30	Repaired & Up	15.5	12363.5		
12/16/89	19:00	Mtr Brg-Shaft End Dn		17059.0	17059.0	Mtr Brg-Shaft End Failure
12/17/89	23:00	Repaired & Up	28.0	17059.0		
4/17/90	5:30	Main Shaft Seal Down		19945.5	7582.0	Main Seal Failure
4/18/90	1:30	Repaired & Up	20.0	19945.5		
12/12/90	15:00	Left Shaft Brg Down		25671.0	25671.0	Bearing Failure
12/12/90	23:00	Repaired & Up	8.0	25671.0		

Table 2: C-A Data

Cum Time (hours)	Cum Failures	MTBF (hrs/failure)
17	1	17.0
21	2	10.5
164	3	54.7
253.5	4	63.4
7902	5	1580.4
12363.5	6	2060.6
17059	7	2437.0
19945.5	8	2493.2
25671	9	2852.3

Consider the data from Table 1 to prepare Table 2 for a C-A plot. Notice the data from Table 1 clearly contains mixed failure modes. Common sense says it is also clear the failures of well designed bearings in Table 1 have lives that are too short as shown by the short mean time between failure (MTBF) for the system.

Inspection of the data that following the coupling failure at time 253.5 hours, shows corrections were made to the system that substantially lengthened the time to the next failure—these conclusions can be obtained simply by inspection.

Table 2 data is used to construct a Crow-AMSAA plot in Figure 1 for the aggregate data. Three lines on the plot show a trend line for all the data and for the two different régimes. You are allowed to use good engineering judgment to decide which trend line best represents the data. The first four data points reflect the poor performance of the system whereby the bearings were killed (most likely because of faulty alignment)—but corrective action has cleared up the infant mortality problems. The last five data points better represent the system response whereas the trend line from using all the data is a smeared-over trend line that is substantially influenced by the bad behavior of early failures.

Figure 1: Crow-AMSAA Plot Of Data From Table 2

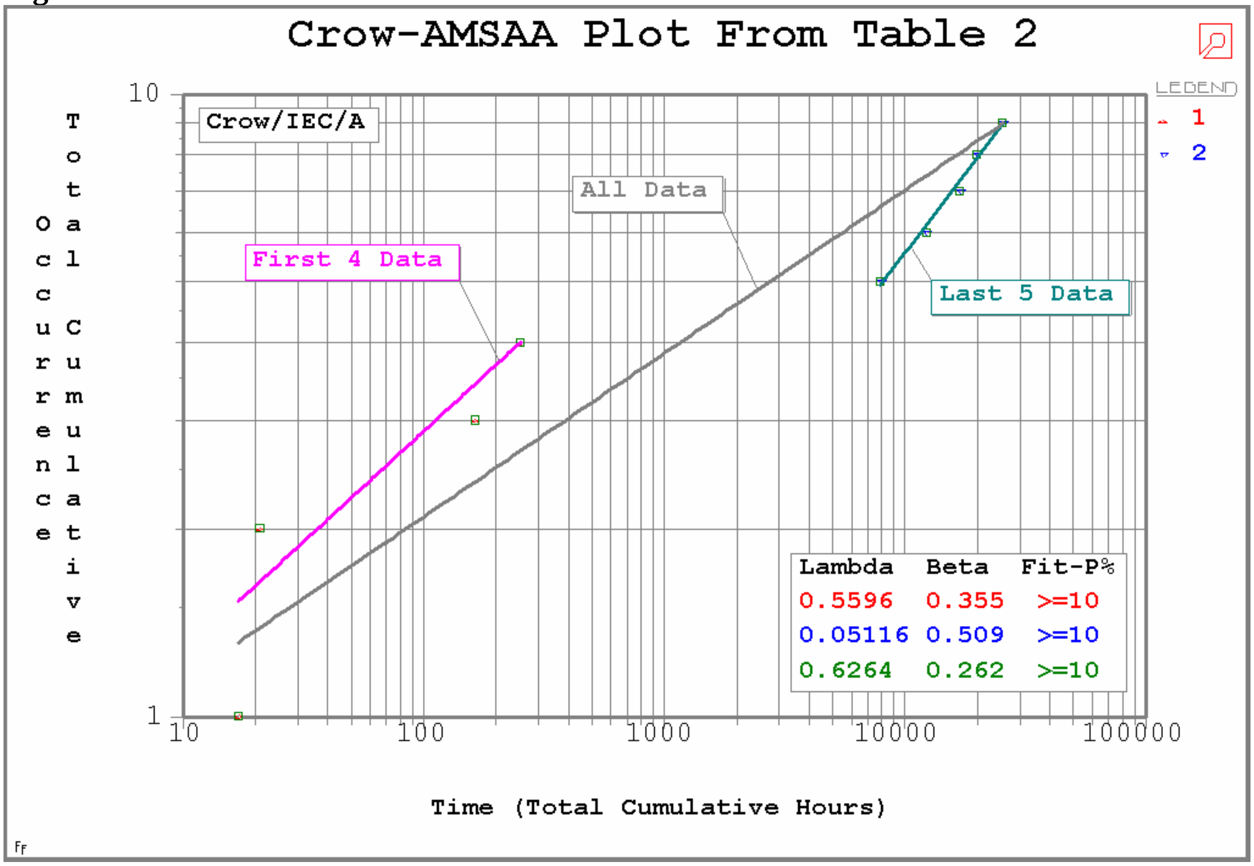

The early failures are represented by the statistics λ=0.5596 and β=0.355 for the "ancient" history. We have little interest in this phase of the data as corrective actions have been implemented.

The most current history is represented by the statistics λ=0.05116 and β=0.509 which will be used to make a "fearless forecast" of future failures for the system using the equation $N(t)=\lambda*t^{\beta}$. Solving the equation for cumulative time gives $t=(N(t)/\lambda)^{(1/\beta)}$. We can forecast the next failure, failure number 10, will occur at $t=(10/0.05116)^{(1/0.509)}$ = 31704.3 hours which is forecasted as 31704.3 – 25671 = 6033.3 hours into the future.

Skeptics will now say you cannot forecast the future with any accuracy. So simply try this experiment using the data in Table 2 as unfolded in Table 3 as the dataset develops over time. The forecasted failures are roughly within 10% of the actual cumulative times as observed from the error data—more data

Table 3: C-A Forecast Using Small Data Set

	Cum Time (hours)	Cum Failures	Cum Time (hours)	Cum Failures	Cum Time (hours)	Cum Failures
	12363.5	6	12363.5	6	12363.5	6
	17059	7	17059	7	17059	7
			19945.5	8	19945.5	8
					25671	9
λ =	0.06587		0.01836		0.05116	
β=	0.479		0.614		0.509	
Fcst (hrs) =	22468.3	8	24083.0	9	31704.3	10
Δ Time (hrs)=	5409.3		4137.5		6033.3	
Error (hrs)=	2522.8		-1588.0		?	

reduces forecast errors. Of course we cannot know the actual time for failure number 10 because it is not listed within Table 1.

So how would you use the fearless forecast of future failures? Use the data for resource planning, develop strategies for preventing the next failure, establish a predictive maintenance watch on the equipment to shut it down just before destructive events occur, etc. Use the data with fearless forecast to also demonstrate your corrective actions have avoided the predicted failures for superior results by avoiding failures. Use the data to avoid the perpetual state of cluelessness so prevalent in many production and maintenance departments! Ah Ha!---If you can predict the failures can you prevent them? Consider the C-A plot for chemical plant safety data shown in Figure 2—then you decide.

Figure 2: Crow-AMSAA Plot Of Actual Safety Data From A Major Chemical Plant

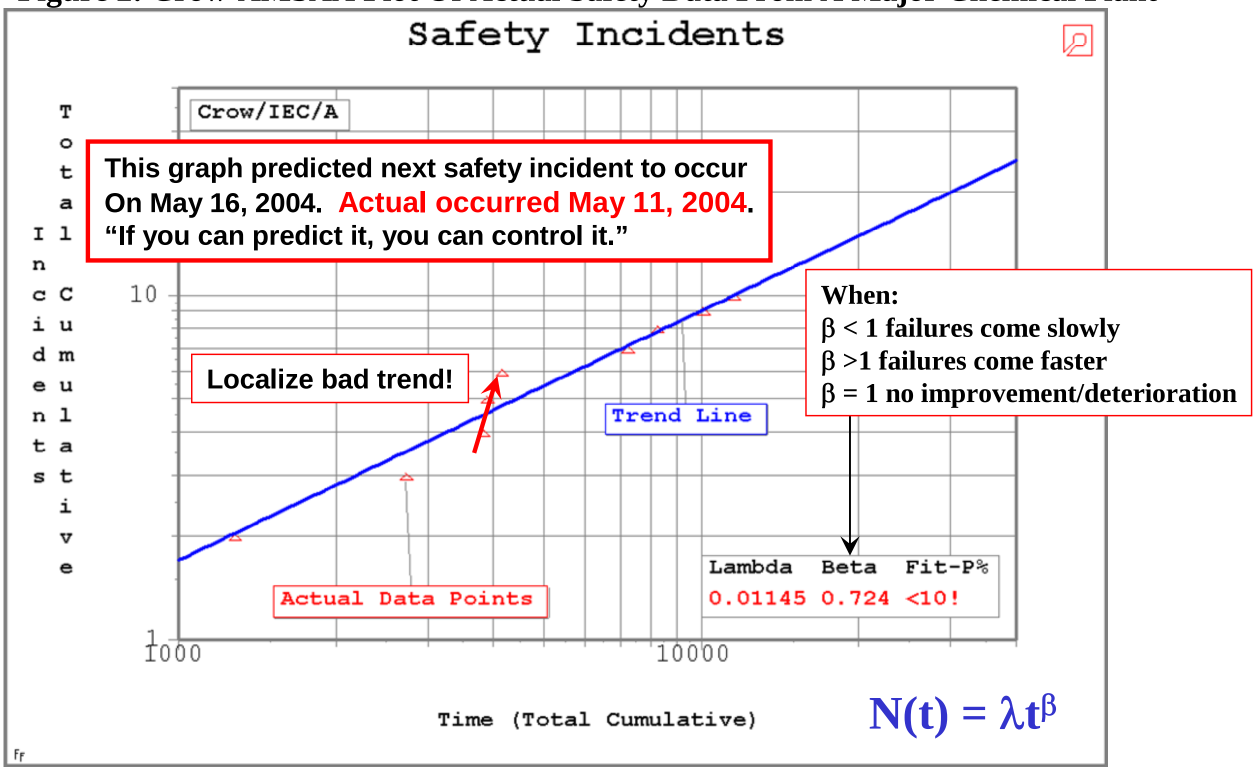

The data for Figure 2 concerning safety incidents is described in Table 4. Clearly safety incidents stem from a multitude of reasons, and clearly every safety incident is a failure of the system to prevent the failure. Most people reject the idea that safety incidents are predictable! However, I have been studying the safety phenomena and straight lines on log-log plots since 1967. Time after time straight line segments appear on log-log paper.

Table 4: C-A Safety Data

Cum Time (hours)	Cum Incidents
1278.9	2
2715.3	3
3813.8	4
3891.4	5
4157.0	6
7223.0	7
8242.0	8
10108.3	9
11560.0	10

When significant improvements occur in safety programs, cusps appear in the data and the trend line forms at a flatter slope to illustrate improvements have been achieved and reductions in failure rates are occurring. Since humans are involved in the safety incidents, you must involve the same humans in the solution and this requires you have simple graphical plots for predicting future events and for displaying significant improvements have been demonstrated. Thus C-A plots are excellent tools for displaying improvements involving the reliability of humans and systems.

Figure 2 also emphasizes the importance of the statistic beta which describes the slope of the trend line. Time after time we see safety programs with cusps forming where the trend line shoots off in the wrong direction with increasing beta slopes and no one is actively working to turn around the undesirable results. If you don't have the clear graphics of a C-A plot you cannot formulate clear battle plans for corrective action. You can conclude the obvious: **You must make reliability issues visible** so each team member can understand what is happening, and make it as obvious as the nose on your face, otherwise, bad results occur and become the accepted norm.

We have other data in the public domain that can help us make predictions concerning future failures. Consider that data from the USA space shuttle. Two shuttles have been lost: 1) The Challenger space shuttle was lost on flight number 25, and 2) The Columbia space shuttle was lost on flight 113. Using these two data points obtained from the NASA database at http://science.ksc.nasa.gov/shuttle/missions/, when is the next space shuttle loss predicted?

Figure 3: Crow-AMSAA Forecast Of Next Space Shuttle Loss

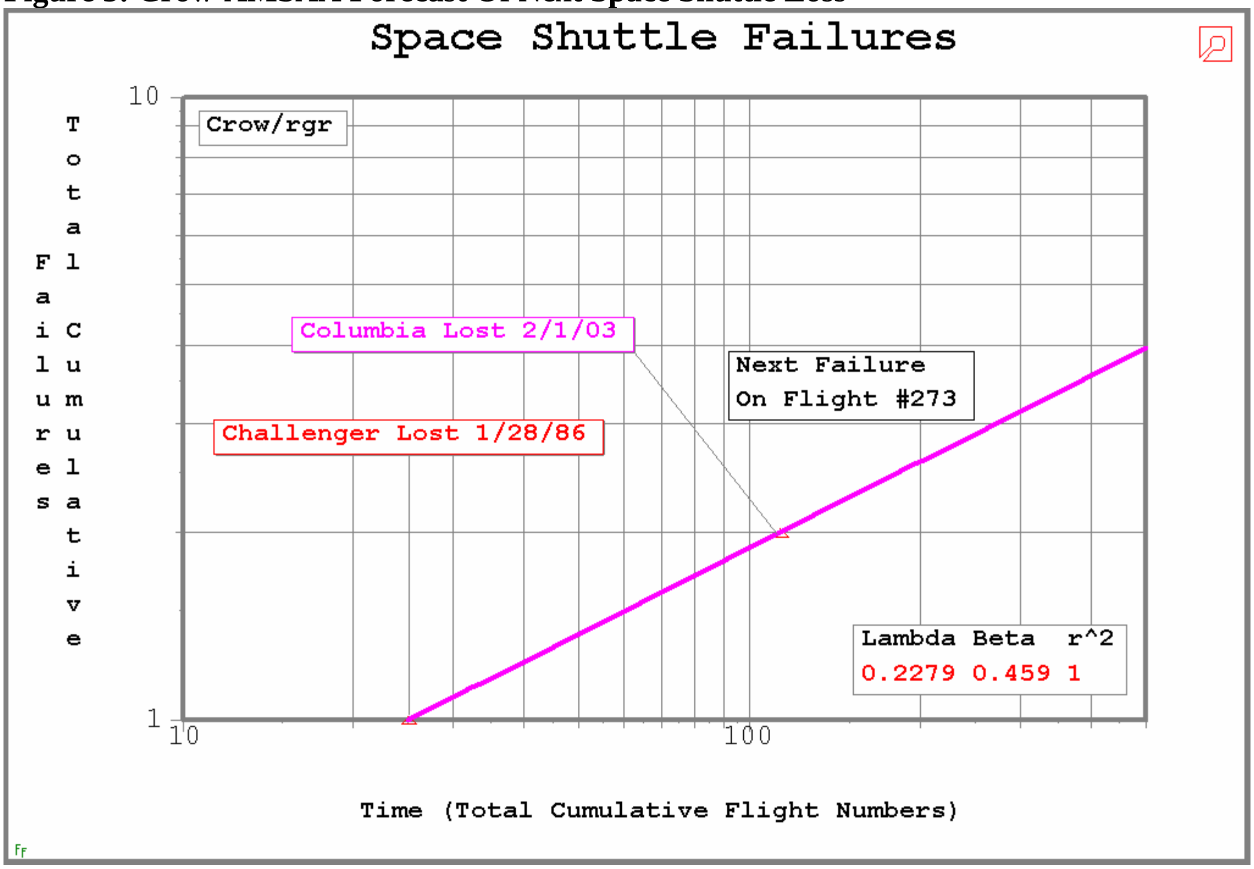

Figure 3 forecasts the next shuttle failure at flight number 273 based on the extrapolated trend line where the line crosses failure #3. Figure 3's line slope shows beta = 0.459 which tells of significant improvements at NASA to grow the time interval between failures even with aging equipment. As noted in Figure 2 when $\beta<1$, failures come more slowly, when $\beta\sim1$, failure rate is unchanged, and when $\beta>1$, failures come more quickly thus the line slope has important physical relationships about failure rates. Both beta and lambda are required for making future forecasts. Your first thought might be how can you predict space shuttle failures based on only two data points: 1) two data points is all you have and two more than you really want to have recorded considering the large loss of life and extreme expense, 2) it is better to make the most reasonable forecast you can to give the data a voice, and 3) the line slope tells if failures are coming slower, faster, or without significant changes. C-A plots are reliability growth plots.

The task of reliability engineers is to make improvements to reduce the failure trend lines and thus make a cusp appear on the C-A plots to signify progress. A clear cusp, based on improvements, are shown for the world's most unsafe railroad (length of the light-rail road is 7.5 miles long) in Houston, Texas as shown in Figure 4.

Figure 4: Houston, Texas Light Rail Line Accidents

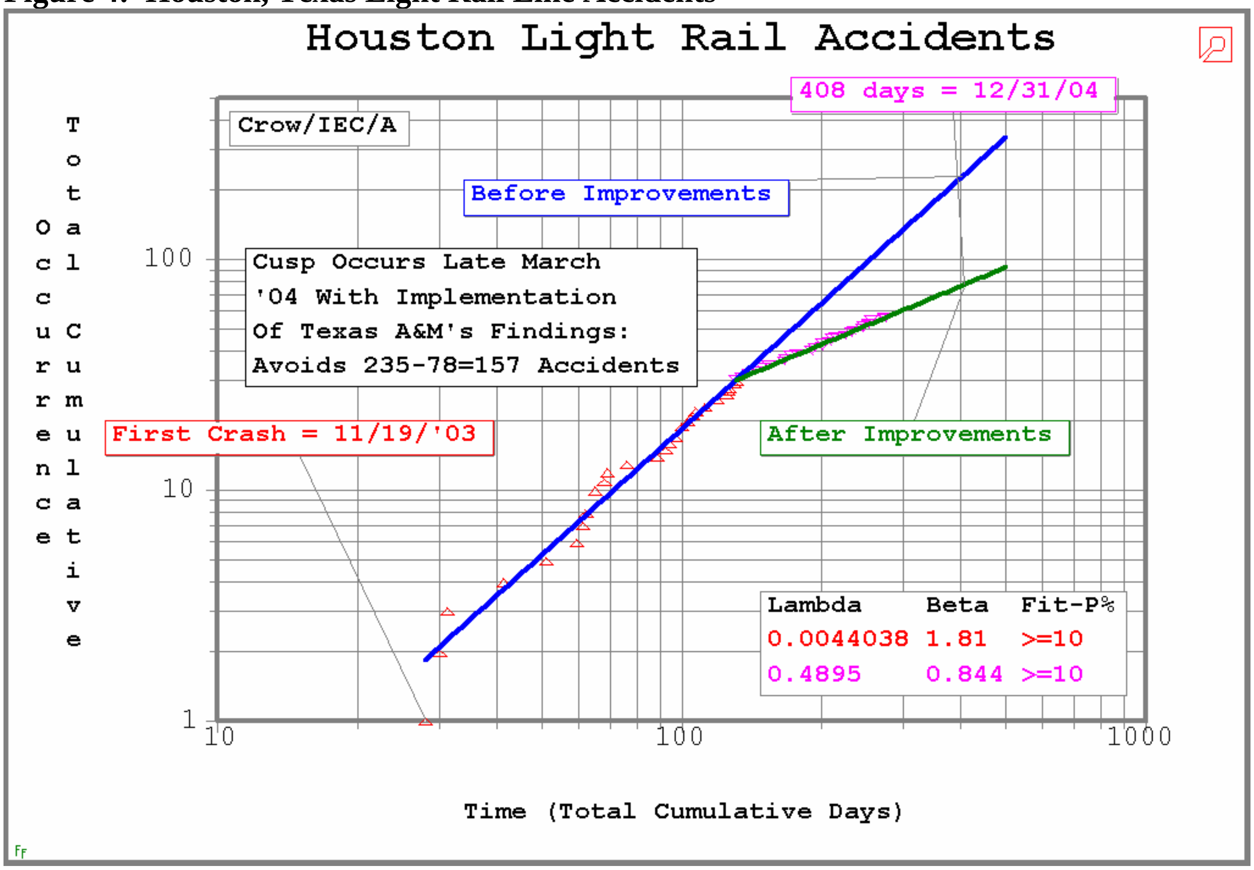

The triangular data points show the early accidents prior to implementing improvements. The inverted triangular data points show a reduction in failure rates after the improvement program. The cusp is clear, and improvement has occurred. Extrapolating the trend line to the end of 2004 suggests 235 failures would have been expected by December 31, 2004 which is 408 cumulative days. Extrapolating the improvement trend line in a similar manner shows 78 accidents are expected. Improvements have occurred, a reduction of 157 accidents are projected to occur in the time interval. The plots are “show me, don’t tell me” about improvements where mixed failure modes occur.

C-A plots similar to Figure 4 visually shows improvements have occurred to reduce accidents. Furthermore the simple graphic allows quantification of how much of an improvement has really occurred during a time interval to justify the improvement programs. Likewise, the line slopes of the C-A plots are helpful in determining how much improvement/deterioration has really occurred. Remember this methodology provides simple forecast from straight lines which makes “sales” presentations easy to understand.

These few examples show the power of simple data sets converted into simple plots to make practical forecast of future failures. More details on Crow-AMSAA plots are available on the Internet (Barringer 2003) and the WinSMITH Visual software is used to make the plots (Fulton 2004).

Weibull Plots 101-

Weibull plots are simple probability plots which handled tailed data of age-to-failure information. Age-to-failure data is plotted in rank order on the X-axis. Median ranks plotting positions are used on the Y-axis to get straight lines on the probability plot. Probability plots usually show the cumulative percentage of the population expected to fail by a given age based on the sample of data supplied for analysis. Weibull plots have been found extremely utilitarian for use with tailed data as usually occurs with the life of components in industry. Thus frugal data sets can provide maximum information without the usual hedge of “assuming the data is reasonably normally distributed” which frequently results from demanding use of Gaussian distributions. The Weibull distribution is the tool of choice for reliability analysis of life data and more information is given in Abernethy (2002).

Weibull plots of component failures by a single failure mode produce three statistics:

1) Line slope beta is often referred to as a shape factor,
2) Location parameter eta which occurs where the trend line cuts the 63.2% CDF (cumulative distribution function) as a mathematical property of the Weibull distribution,
3) Goodness of fit parameter r^2 known as the coefficient of determination or a pve% know as a p-value estimate.

For component Weibull plots of single failure modes, the Weibull line slopes, beta (not to be confused with the C-A line slopes also known as beta), have physical significance:

1) Beta < 1, infant mortality characterized by a declining instantaneous failure rate with time,
2) Beta = 1, chance failures have a constant instantaneous failure rate with time, and
3) Beta > 1, wear out failures characterized by increasing instantaneous failure rate with time.

The important issue is that Weibull plots of component failure data tells you **how** things die. A typical observation by non-Weibull analysis often defines a failure as “wear out” whereas performing a Weibull analysis of component failures provides considerable enlightenment that the motivator for end of life is a different failure mechanism than verbalized.

For selection of a maintenance strategy it is very important to know the actual failure mode for components characterized by the Weibull beta value:

1) Beta < 1, for infant mortality, is a run to failure strategy. An old part is better than a new part because the failure rate is lower as weak units have been eliminated from the population.
2) Beta = 1, for chance failures, is a run to failure strategy. An old part has the same failure rate as a new part. Thus nothing is gained by a replacement strategy which throws away unused life until the failure mode changes to a wear out mode.
3) Beta > 1, for a wear out failure, <u>may</u> have an optimum replacement strategy <u>if</u> the cost or safety consequences have a very high cost ratio for an unplanned failure compared to a planned replacement cost which then drives a preventive replacement strategy. if unplanned replacement cost or safety costs are not much greater than a planned replacement costs, then the maintenance replacement strategy may be run to failure controlled by economics.

The basic issues are to know the signals for taking action and tempering the signals by the economics. Weibull beta values are driven by the physical failure data taken from CMMS. In short, let the data talk and give it a voice through Weibull analysis.

Consider the data in Table 5 containing age-to-failure data for a rotating seal. The data produces the Weibull plot in Figure 5. The age to failure data is based on the operating ages derived from the birth date and the death date—in short, how old was the component before it failed. The Y-axis plot position information is obtained from Bernard's median rank equation. Bernard's plot position is generally considered best practice for reliability data and is explained by the equation (i-0.3)/(N+0.4) where i is the rank position for ordering the data and N is the number of data points in the data set. Bernard's equation is explained in Abernethy (2002).

Table 5

Age-To-Failure (days)	Order	X-axis: Age-To-Failure Ranked	Y-axis: Plot Position %
1264	1	25	7.45%
1339	2	163	18.09%
25	3	224	28.72%
359	4	244	39.36%
537	5	359	50.00%
224	6	537	60.64%
163	7	775	71.28%
244	8	1264	81.91%
775	9	1339	92.55%

The Weibull plot in Figure 5 shows a good curve fit with the coefficient of determination (r^2) greater than the critical value of 0.8464 which leads us to conclude that all is well with the Weibull plot. However, you should consider why a well designed seal, properly operated, should have an infant mortality failure mode (β<1) rather than a wear out failure mode which was expected!

Figure 5: Weibull Plot Showing Infantile Mortality

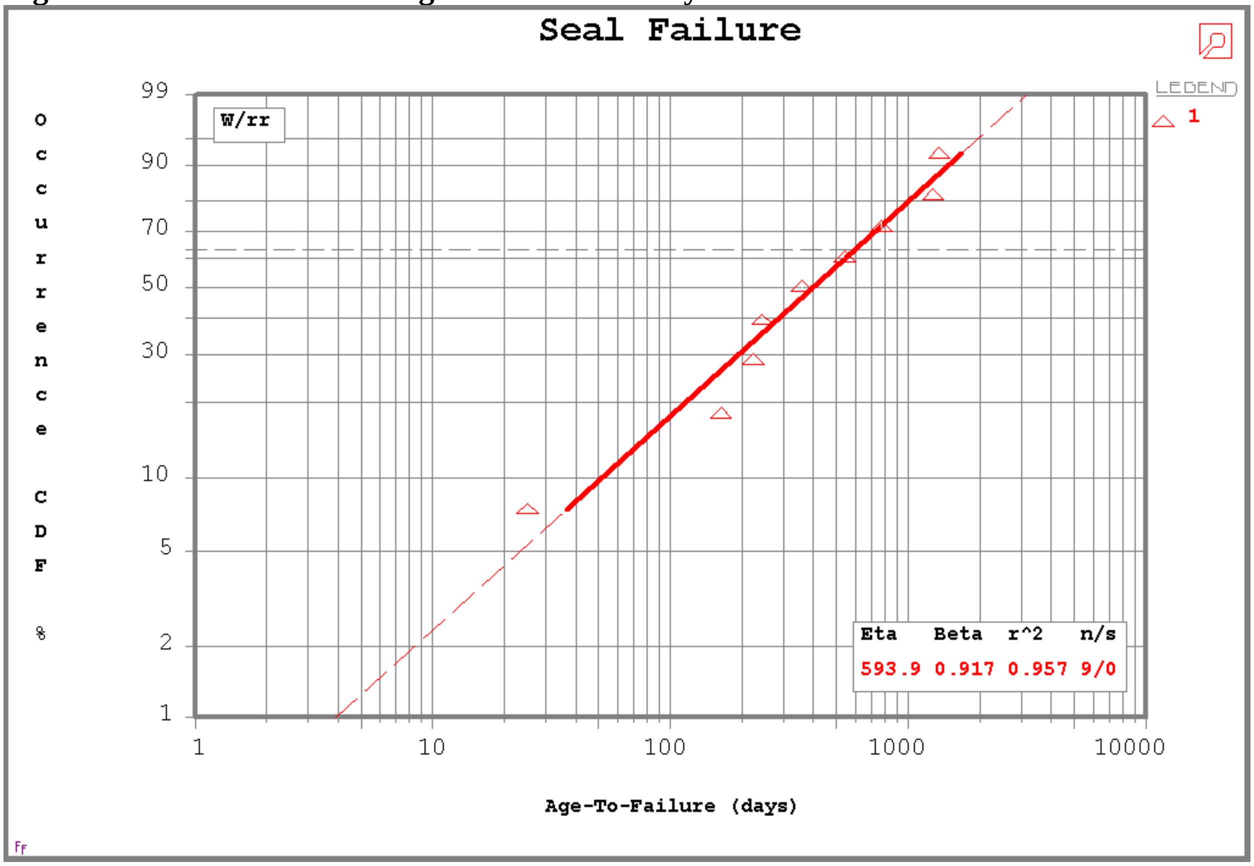

Further consideration of the data shows that some failures resulted from process errors that killed the seal rather than the seal dying on its own. Thus we must handle the data in censored (suspended) form rather than naively taking the data at face value based on a good curve fit! Basically, Figure 5 tells you how the <u>system</u> is performing rather than how the <u>seal</u> is performing because the system has mixed failure modes.

Table 6 shows the corrected data. Bernard's median rank is corrected for the suspended data, shown with a minus sign to signify it is suspended. A careful study of the data resulted in ~63% of the data representing suspended data which tells you immediately that to achieve long seal life you must correct the genocide occurring from how the seal is used. The correction in Bernard's equation is achieved by use of Auth's approximation which is also explained in Abernethy (2002). For most of us, it is impossible to remember all the arcane rules associated with how to calculate the plotting position, so it is an advantage to use WinSMITH Weibull software (Fulton 2004) which follows all the best practice rules outlined in The New Weibull Handbook. Please note that suspended data is not plotted but the suspensions are used in the data set for corrected plotting positions.

Table 6

		Corrected For Suspensions	
Age-To-Failure (days)	Order	X-axis: Age-To-Failure Ranked	Y-axis: Plot Position %
1264	1	-25	suspended
1339	2	-163	suspended
-25	3	-224	suspended
-359	4	-244	suspended
537	5	-359	suspended
-224	6	537	18.09%
-163	7	-775	suspended
-244	8	1264	46.45%
-775	9	1339	74.82%

Figure 6 shows the plot of data from Table 6.

Figure 6: Seal Failure Data With Suspensions Correctly Identified

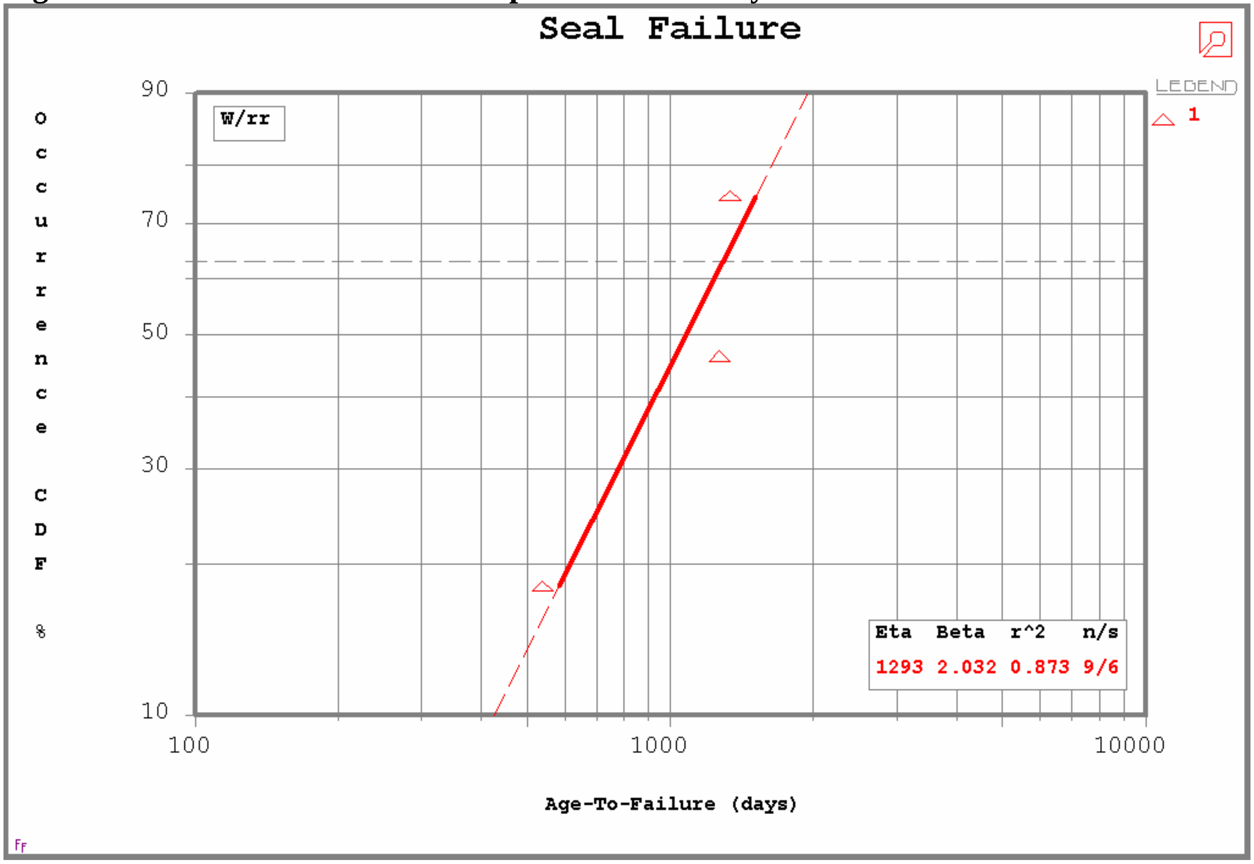

Notice the β>1 tells you the seal has a wear out failure mode and the r^2 value represents a good curve fit for the three ages to failure when considered against the critical value of r^2=0.7921. The eta value says the seal has a characteristic life of 1293 days when it is operated without genocide conditions. This clearly says we need to correct operating conditions to get long seal life and a wear out failure mode.

Maintaining a Weibull failure data base for your operating conditions provides important history. For example, many heat exchangers have a wear out failure mode with $3 < \beta < 15$ depending upon the specific way the tubes can fail. Usually we can expect to see a general corrosion $\beta \cong 3$, however, if the corrosion rate is severe, then $\beta \cong 9$, and if bad things are happening very fast, then $\beta \cong 15$.

Consider the heat exchanger tube life data in Table 7. When will the next heat exchanger tube fail? The data represents one failure at age 7 years with 142 survivors at age 7 years. All future failures will come from today's survivors. Failed tubes are usually plugged rather than replaced with roughly 10% loss of tubes representing the point where the heat exchanger has reached end of life.

Table 7

Age-To-Failure (years)
7
-7*142

Using the range of beta values from the Weibull data base we can make three Weibayes forecast values: one forecast for β=3 will show the next failure occurring far out into the future, one forecast for β=9 which will show a mid range forecast, and one forecast for β=15 will give the most pessimistic time for when the next failure will occur. Taking the data from Table 7 and making a Weibayes forecast using the beta value ranges with WinSMITH Weibull software will allow an Abernethy risk forecast which is summarized in Table 8. Table 8 tells us what to expect into the future. We can set up a watch to test for future failures. Roughly 80% of heat exchangers will display a β≅3 so we have good odds that a calamity might not occur. Forewarned is forearmed for prudent management.

Table 8

	Years Into The Future For Forecasted Failures			
Age-To-Failure (years)	**Future Failures**	**β=3**	**β=9**	**β=15**
7	1st failure	1.8	0.6	0.3
-7*142	2nd failure	3.1	1.0	0.5
	3rd failure	4.2	1.2	0.8

Summary

Crow-AMSAA reliability growth plots with straight lines on log-log plots make reliability data visible. Crow-AMSAA plots allow demonstration of failures avoided by corrective action techniques. Crow-AMSAA plots work well with mixed failure modes and will provide useful information without starting the plots at time zero. Crow-AMSAA plots allow forecast of future failures for hardware and safety events.

Weibull plots of component failures using age-to-failure data provide useful information about failure modes with the trend line statistic beta. Knowing the beta value allows structuring the maintenance strategy. Useful libraries of information are obtained from knowing the line slope beta values and characteristic age to failure eta statistics. Knowing how things fail gives detailed forecast of future failures which is important for logistics and maintenance manpower demands.

Make your data talk. Provide the graphics to help sell your ideas. Prevent failures and save money.

References

Abernethy, Robert B, **The New Weibull Handbook**, 4th edition, published by the author, North Palm Beach, FL, 2002

Barringer, H. Paul, **November 2003 Problem Of The Month—Crow-AMSAA Plots**, http://www.barringer1.com/nov03prb.htm , 2003.

Fulton, Wes, **WinSMITH Visual** software, version 4.0WJ, Fulton Findings, Torrance, CA, (2004)

Fulton, Wes, **WinSMITH Weibull** software, version 4.0WJ, Fulton Findings, Torrance, CA, (2004)