TROUBLESHOOTING TURBOMACHINERY PROBLEMS USING A STATISTICAL ANALYSIS OF FAILURE DATA

by

James E. Corley

Senior Engineering Consultant

Saudi Aramco

DHAHRAN, SAUDI ARABIA

James E. Corley is with Saudi Aramco in Dhahran, where he has corporate responsibility for technical matters relating to dynamic analysis and rotating equipment. He currently holds the title of Senior Engineering Consultant which is the highest technical position in Aramco. He joined Aramco in 1973 to organize the company's vibration analysis unit in the Consulting Services Department. He has also developed and implemented a machinery monitoring program which is used throughout the company. He is currently developing a program to improve the reliability of pump mechanical seals.

Prior to his present employment, he was with Union Carbide Corporation's Nuclear Division in Oak Ridge, Tennessee. While there, he was responsible for compressor seal development for the gaseous diffusion complex. Mr. Corley holds B.S. and M.S. degrees in Mechanical Engineering from Mississippi State University, and has done post graduate work at the University of Tennessee.

ABSTRACT

The use of statistical techniques to investigate machinery failure data is one of the most powerful troubleshooting tools available to the turbomachinery engineer. This analysis tool can provide a great deal of information and insight concerning causes of failures and the probability of future ones. Using a graphical technique known as a hazard analysis or Weibull analysis, failure data of the type commonly available in the field can be analyzed to determine failure modes, the mean time to failure, the effects of different failure modes and the probability of future failure.

The graphical procedure described in this paper is very simple to use, and requires only elementary math. The technique uses such readily available parameters as running hr or installed time of the subject under investigation and can easily handle incomplete or censored data.

The method used to construct and interpret a hazard, or Weibull chart, using either special Weibull paper or simple log-log graph paper is presented. A simple fictitious example is used to illustrate the method. Eight case histories are then presented to illustrate the procedure and show the many ways that the analysis can be applied to machinery problems.

INTRODUCTION

For anyone who has been asked "What are the chances of this machine lasting until the next turn-around?," or "What is the average life of the mechanical seals in the plant?," or "How much good did the modification do to improve the life of a machine?," the use of the statistical analysis of failure data called a hazard analysis can provide answers. This technique is a powerful tool for the machinery engineer who often must try to make sense of a bewildering amount of data that is often incomplete. In other words, "Real field data."

The type of failure data that is normally available to the machinery engineer generally consists of both information on failed machines and on machines that are either still operating, or have been removed from service for reasons that are unrelated to the failures. For example, a seal under study that has not failed may be changed out during a pump overhaul. The information of interest may be the running hr, the time in service, or possibly the number of starts on the machine. The statistical technique used to analyze this type of data is called a hazard analysis of censored data. Censored data consists of failure times or running time (in any convenient units) which are mixed with the running times of unfailed machines. A special case of the hazard analysis is also known as a Weibull analysis.

Although computer programs are available for analyzing large amounts of data, say more than 100 points, most analysis requirements can be easily met using only a calculator and graph paper. Data plotted on these charts can give the engineer insight into the type of failure modes that are predominant in the data, such as a wear in or wearout mode or a constant failure rate mode. This technique is also used in determining if a failure mode is correlated to running time or possibly another parameter such as number of starts. Because the plotted hazard curve is related to the more conventional "bath tub curve" of failure rate as a function of time, it can also be used to estimate the time to wearout. This is useful in determining preventative maintenance intervals.

Normally in an investigation, there are several machines or components that have failed and many that are still running. The information pertaining to these "nonfailures" must also be considered in any analysis, and the hazard analysis automatically takes this into consideration.

Although the shape of the hazard curve is frequently of interest in itself, an accurate knowledge of the life is of particular importance in assessing the effects that design changes may have. The hazard analysis allows one to separate out the effects on life of different or competing failure modes. If a component such as a mechanical seal has multiple ways in which it can fail, one can estimate the life improvement to the population if one of these modes is eliminated by a design change. The analysis can also be used to measure the effects of that design change when it is implemented.

Hazard analysis also provides an estimate of probability of failure. In a population of machines for which data is available, it is possible to calculate the probability of failure in a future time span, such as the period until the next plant shutdown. This gives plant management the quantitative information with which to make decisions.

HAZARD ANALYSIS

The hazard function can be thought of as the instantaneous failure rate of a set of failure data. It is a measure of the proneness to failure of a machine as a function of age or running hr. The Weibull function is a two parameter probability function that describes a time dependent failure rate. The Weibull cumulative distribution function is given by:

$$F(x)=1-e^{-(x/\alpha)\beta} \tag{1}$$

where α is the value of the scale parameter, which is a characteristic time to failure, and β is the value for the shape parameter.

Although there are many different types of probability functions that may work or fit the data, such as an exponential, normal, log normal, or extreme value distribution, the Weibull function seems to fit almost all of the data that is commonly found with machinery failures. Because of this generality, the Weibull analysis will be used exclusively herein. Indeed the Weibull distribution incorporates some of the above distributions, such as the exponential and normal as special cases. Herein, the terms "hazard analysis" and "Weibull analysis" may be used interchangeably.

Although this function is, in general, a complicated mathematical probability distribution, as seen above, this discussion will leave the derivations and proofs of the higher statistical mathematics to others. For those that are interested, the full math treatment is given by Nelson [1].

The math used herein will be no more than simple arithmetic and a simple graphical procedure. This graphical method was developed by Nelson [2] and Bianco [3], and is described in detail in two excellent papers by Nelson [4,5]. This procedure uses either special graph paper [6] or, in the case of the Weibull analysis, simple log-log graph paper.

Bath Tub Curve

One of the most powerful aspects of the hazard analysis is the insight that it can yield into the mechanism of failure modes. Most people are familiar with the concept of the "Bath Tub Curve" to describe the change of failure rate with time. Many mechanical and electronic devices have a typical failure rate curve which is a function of age. This characteristic is shown in Figure 1, and the shape of the curve is responsible for the name.

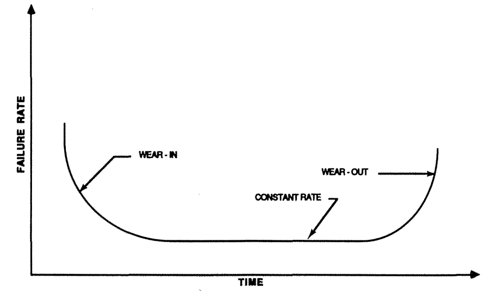

Figure 1. Bathtub Curve Showing Failure Rate as a Function of Time.

At the beginning of its life, a device may exhibit a high failure rate that decreases with age or time. This charactistic can be due to problems of design, manufacture, assembly, or installation. The failures during this period are commonly known as "infant mortality" or wear in failures. Many electronic components display this behavior and may be "burnt in" before being put into service to eliminate this type of failure mode.

As weak components fail, the curve levels out to a failure rate which remains constant for a period of time. In this mode, the constant failure rate produces an exponential probability distribution. During this period, the failure rate is independent of time in service and the probability of failure in a given time increment is the same regardless of the age of the component. For example, if a mechanical seal displays this type of distribution, the probability of its failing in a six months period is the same, regardless of whether it has been in service for only a few weeks or for several years.

When a mechanical component, such as a seal or antifriction bearing, reaches the end of its useful or design life, the failure rate will start to increase with time. This is the period of wearout failure and is shown as the backend of the bathtub. Devices that are prone to wear or fatigue will generally display this type of behavior.

Weibull Data Plotting

While the "Bath Tub" is a useful concept, in practice it may be impractical to plot due to lack of sufficient failure points to define the curve and complicating factors such as censored data (nonfailures). A hazard plot on Weibull paper (log-log) will display the same characteristics as that of the "Bath Tub curve" and will yield several qualitative results that are extremely valuable.

A generalized hazard plot on Weibull paper (shown in Figure 2) has the same type of information as the "bath tub" except that it is flipped over and rotated. To construct this plot, failure data are plotted on the ordinate as a function of a parameter called the cumulative hazard. The method of calculation of the cumulative hazard will be explained later but it has no real physical interpretation. The ordinate used might be in some convenient units of time, such as running hr, installed hr, or miles traveled. In some cases, as in electric motors or combustion gas turbines, the parameter of interest might be the number of starts on the unit. The units of measure used are set by the particular problem under consideration and are determined by the engineer.

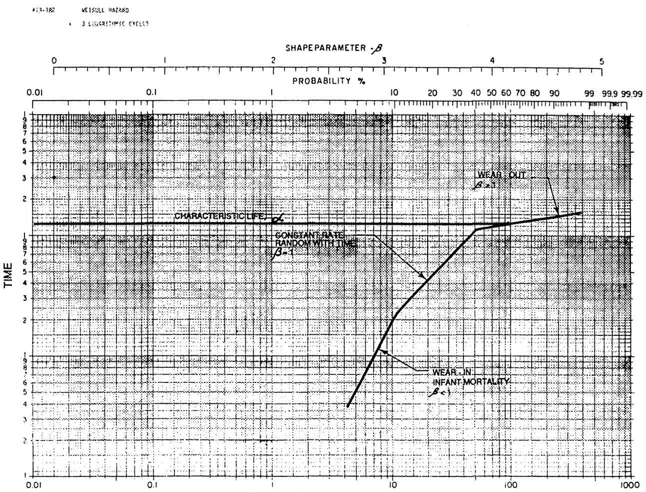

Figure 2. Generalized Hazard Plot of Typical Data.

On this Weibull hazard plot, the infant mortality portion is represented by a line which has a slope that is greater than 1.0. In this region the failure rate is decreasing with time. If the failure data plots out in this manner, one should consider that the problem might be one of assembly or installation. The longer it runs the less likely it is to fail.

The constant failure portion of the "bath tub" is shown on the Weibull plot as a line with a slope that is equal to 1.0 in Figure 2. Data that falls on this section of the curve have an exponential type of failure distribution and are generated by mechanisms that are independent of time in service. An example of this would be seals that fail because an operator ran a pump dry; a situation that might arise which is independent of when the seal was installed.

Failures that are caused by wearout type of mechanisms will plot out as a line with a slope less than 1.0. In this region, the failure rate is increasing with time and the longer a component is in service, the greater is the probability that it will fail. Ball bearings with their limited life due to fatigue would display this characteristic, as would mechanical seals that fail due to normal wear. If it hasn't been noticed, people display this same sort of "bath tub" failure distribution.

Although some machines may display all of the above characteristics, from "wear in" to "wearout," in most cases of failure analysis there is only a single failure mode of interest and the data will show predominately only one single line when plotted. The full range of the "bath tub" will be evident when there are several different failure modes. When the data is plotted, as will be explained later, a great deal of insight into the type of failure under investigation can be gained just by looking at the shape and slope of the curve. If, for example, the curve has a slope > than 1.0, which indicates an "infant mortality" type of failure, one should look closely at installation or assembly problems. If, on the other hand, the plotted failure data results in a line with a slope that is < than 1.0, the failure mechanism will be one that yields a wearout type of mode, such as fatigue or corrosion.

Weibull Shape Parameter, β

The slope of the curve on the Weibull plot is the inverse of the Weibull shape parameter β that was defined in Equation 1. An example, using fictitious data for illustrative purposes, is shown in Figure 3. In this plot, β is determined by drawing a line parallel to the data line that passes through the dot located in the upper left side of the graph paper. This line is then extended to intersect the Shape Parameter scale at the top of the figure and β is read off directly. In the example shown in Figure 3, β is 1.8. If simple log-log paper is used, the slope of the data line can be measured directly and in this example is 0.55. The shape parameter β is simply the inverse of this slope. Since β is the inverse of slope, it follows that for a wear in type of failure mode, β is > than 1.0. For a constant failure rate, β is equal to 1.0, and for a wearout failure mode, β is < than 1.0.

Characteristic Life and Mean Time Between Failure (MTBF)

In investigating many failure analysis problems, a parameter of great interest is the average life of the population. In a hazard analysis, the average or characteristic life can be read off directly from the data plot as the value of time where the curve crosses the cumulative hazard value of 100. From Equation 1, this characteristic life is the Weibull parameter, alpha. In the general case, the Mean Time Between Failure (MTBF) is a function of alpha and β and is given by:

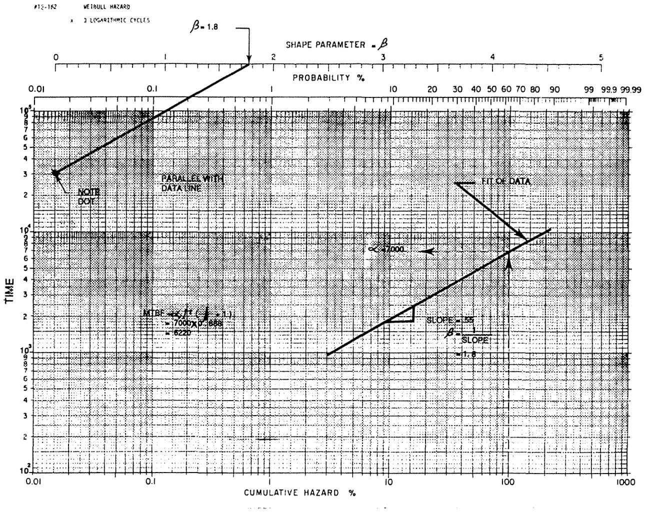

Figure 3. Graphical Determination of the Weibull Parameters α and β Using Weibull Hazard Paper.

$$\text{MTBF}=\alpha \cdot g \cdot [1+(1/\beta)] \qquad (2)$$

where γ is the gamma function. For the special case of the exponential failure distribution, where β is equal to 1.0, the MTBF is equal to the characteristic life, alpha. For comparative purposes, say when one is wanting to determine the effects of a design change on the life of a component, α can be used directly as the effective life without having to resort to the more complicated gamma function. This has the practical advantage of being able to be read directly from the plot.

In the example shown in Figure 1, α is seen to be 7000. With a β of 1.8, the MTBF is calculated, using Equation 2, to be 6220. Thus, for this particular set of data, the failure mode is a wearout mode with an average life of 6220 hr.

Probability of Failure

For a given set of failure data, a plot on Weibull hazard paper can be used to determine the probability of failure using the probability scale at the top of the chart in Figure 4. To determine the probability of failure in a given time, say 3000 hr; enter the y axis at 3000 hr and move out horizontally on the graph to the point of intersection with the data plot at a cumulative hazard of 23. Then move vertically up and read 20.0 percent from the probability scale. Thus there is a 20 percent probability of a failure of the unit by the time it has reached the age of 3000 hr.

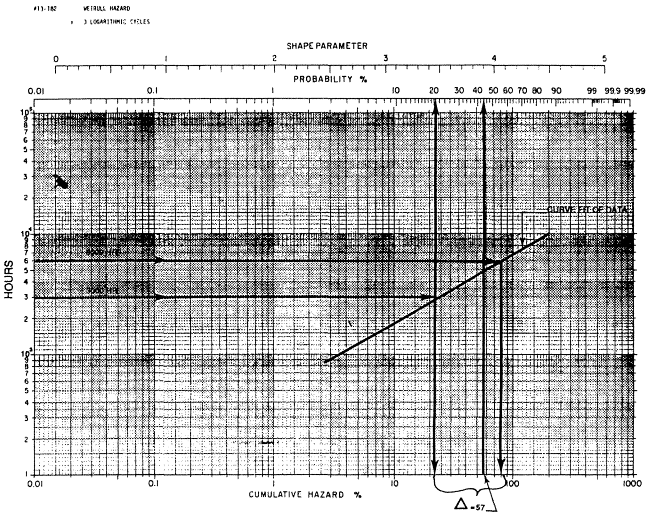

Figure 4. Determination of Probability to Failure Using Weibull Hazard Paper.

Frequently for the machinery engineer, a more common use of this probability function is the determination of chance of failure of a machine in a given time span, say until the next turn around of the plant. In the example shown in Figure 4, say the question was "What is the chance of failure in the next four months (3000 hr) for a machine that already has 3000 running hr?" In this case, the cumulative hazard is determined from the curve at both 3000 and 6000 hr, 23 percent and 80 percent respectively, and the difference in cumulative hazards is found to be 57 percent. The probability to failure is then found by entering the graph at a cumulative hazard of 57 percent and moving vertically upward to the probability scale and reading 44 percent. Thus, the probability of a unit with 3000 running hr failing in the next 3000 hr is 44 percent. Here, an example is seen of the effect of a failure rate that increases with time ($\beta > 1$). There is a 21 percent chance of failure for the first 3000 hr but a 44 percent chance of failure in the next 3000 hr.

In statistical modelling, such as using the Monte Carlo technique, it may be desirable to simulate running time to failure using a random number generator and a statistical distribution. For a given straight line data curve from a Weibull plot, and thus knowing α and β, it is possible to relate probability and time (or other parameter) through the following relationship:

$$T=\alpha \cdot (-\ln (1-P))^{1/\beta} \qquad (3)$$

where P is the probability to failure and T is the time parameter of interest. Conversely, if the time, T, is known, the equation can also be solved for the probability, P.

Different Failure Modes

A common situation in many failure analyses is the presence of different and competing failure modes in the data. For example, an analysis of a mechanical seal population may indicate that the seals fail from several causes, such as buildup of scale, running dry, or an O-ring material failure. The hazard analysis allows these different failure modes to be analyzed and the MTBF determined for each mode. This allows the engineer to evaluate which modes are significant and what improvement could be expected in the overall life if one of the modes is corrected. When a design change is implemented, such as an O-ring materials change, the technique is used to determine the effects on the average life of the population.

In using this technique, label each different failure mode with a unique tag, say "S" for scaling, "O" for O-rings, "P" for overpressure, etc. To determine the MTBF for the O-ring failure mode, all of the other failure data are treated as censored data, that is, as nonfailures, and the hazard analysis is performed using only the O-rings data as failures.

To determine the potential benefits of eliminating a particular failure mode by a design change, as in changing the O-ring material, all of the other modes "S," "P," etc., are considered as failures and the "O" mode is censored. That is, the running times for the O-ring failures are included in the data, but are treated as if they were removed from service without failing. The increase in life or MTBF for the population is the measure of improvement possible by eliminating the O-ring failure. A detailed example of this use of the hazard analysis will be given in Case History 2.

CREATING A HAZARD PLOT

To create a hazard plot, one must first obtain a quantity of failure data. The data for the failed components is generally available in some form, either as running hr or, in the worst case, installed date and failure date. Although installed time is not as precise as actual running hr, it is still a useful measure and is often all that is available. In addition to the time on the failed units, the time on the unfailed units is also needed. This includes not only the time on those that are still running, generally a plant full, but also the time on those units that have been removed from service for reasons other than failure. Remember, these nonfailure times may represent the bulk of the time of the population and must be included.

Example-Hazard Analysis of Widjets

Table 1 will be used as an example for showing how the hazard analysis is performed and plotted. The data shown here has been generated artificially to demonstrate the method only. It represents the running hr for a new design of widjet. The company that manufactures the new widjet, would like to set the warranty period at one half of a year and would like to have no more than one percent of the units fail in that time. The life of the units should be at least six years, and it is expected that the normal failure mechanism would be wearout. To maintain turnover of the product and keep the demand up, it is desirable that no more than one percent of the widjets survive more than 18 years.

Table 1. Example of Hazard Analysis of Widjets Using Simulated Data.

FAILURE DATA MONTHS	ORDERED DATA	REVERSE RANK N	HAZARD 100/N	CUMULATIVE HAZARD
9 F	1 *	18		
1 *	1.5 F	17	5.88	5.9
1.5 F	2 F	16	6.25	12.1
4 F	2 *	15		
7 F	3 F	14	7.14	19.3
4 F	3 *	13		
8 *	4 F	12	8.33	27.6
3 *	4 F	11	9.09	36.7
7 *	5 F	10	10.00	46.7
12 F	5 *	9		
2 F	6 F	8	12.50	59.2
5 F	7 *	7		
8 F	7 F	6	16.67	75.9
2 *	8 *	5		
5 *	8 F	4	25.00	100.9
9 *	9 *	3		
3 F	9 F	2	50.00	150.9
6 F	12 F	1	100.00	250.9

F - FAILURE
* - NONE FAILURE

Having gotten the data and noted which are failures and which are nonfailures, (Column 1 in Table 1) the data is then ranked in ascending order of increasing time (Column 2). This sorting, while not difficult, can be laborious by hand for large sets of data, say more than 50 points. Next, a set of numbers, N, is generated by reverse numbering the data column (3). Note that a value of N is generated for each value of data, including the nonfailure data. A hazard value (Column 4) is then calculated by dividing the number 100 by each value of N that corresponds to a failure. *Important-A Hazard Value is Not Calculated for the Non-Failure Data!* The cumulated hazard (Column 5) is obtained by the cumulative sum of the hazard values in column 4.

A hazard plot is now obtained by plotting the running time of the failed points as a function of the cumulative hazard that was calculated for that point. As mentioned before, the data is plotted on log-log paper or special Weibull paper. The data for widjets is shown plotted in Figure 5. As seen in the figure, the data plots as a reasonably straight line which can be fitted to a straight line either by a least squares curve fit routine or by eye. Using a computer generated curve fit routine allows one to obtain confidence limits as shown in the figure.

For this case of widjet failures, it is seen from the figure that the average life or MTBF is 6.8 years and with a β of 1.74, the failure is concluded to be a wearout mode. The probability of failure in one half of a year, the warranty period, is found by entering the time axis at 0.5 years and going out to the fitted curve. Moving vertical from that point to the probability scale at the top of the figure, the percent of failure is seen to be 0.8 percent. This meets the corporate goal of less than 1 percent.

To find the percent of widjets that fail in 18 years, the above procedure is repeated by entering at 18 years and reading the percentage off of the scale at 99 percent. Thus, the widjet design

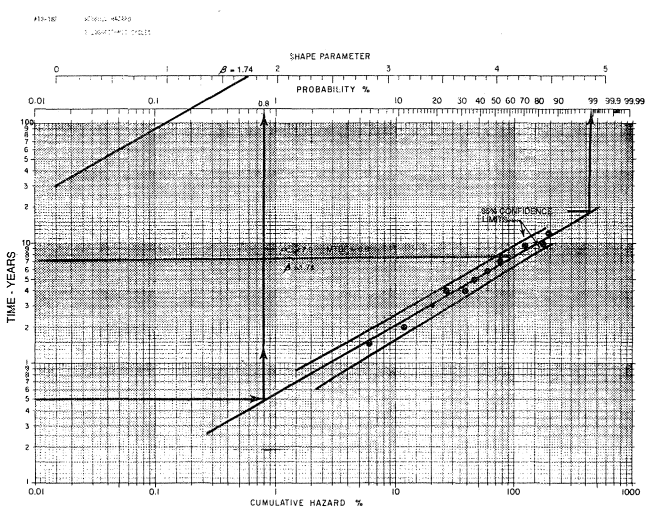

Figure 5. Hazard Plotting of Widjet Failures.

meets all of the corporate requirements within a 95 percent confidence limit.

CASE HISTORIES

Now for the real world. The following case histories are ones that have been found in typical petrochemical facilities and are taken from real problems.

Case History 1—Electric Motor Insulation Failure

This case involves a turn-to-turn insulation problem in a population of 21000 hp, four pole, synchronous electric motors. The motors are used primarily as compressor drives in several gas plants. Approximately two to three years after startup of the plants, the motors started to fail. The mode of failure was breakdown of the turn-to-turn insulation in the coils. Insulation failure of this type is normally caused by high magnitude, very short duration electrical surges, such as lightning (a rather rare occurrence in Saudi Arabia).

A hazard analysis was performed after 7 of 22 machines had failed. This analysis is seen in Table 2, where the parameter under consideration is the running hr of all motors. In plotting the data on Weibull paper in Figure 6, it is seen that the average life is 34651 hr and that beta is 1.13. It should be noted here that the life of the motors is estimated by an extrapolation of the curve, even though none of the machines had attained running hr near the calculated life. Also of interest, because it sheds light on possible failure modes, is the fact that a β of 1.13 indicates that the failure mode is essentially independent of running hr. What ever was causing the failures was not strongly dependent on the operating hr on the motors.

Table 2. Hazard Analysis of Insulation Failure Data of 21000 HP Motor Using Running Hours.

HOURS	REVERSE RANK N	HAZARD 100/N	CUMULATIVE HAZARD	HOURS	REVERSE RANK N	HAZARD 100/N	CUMULATIVE HAZARD
87	28			12023*	13	7.7	19.9
1892*	27	3.7	3.7	12446	12		
2208	26			12795	11		
3503	25			14242*	10	10.0	29.9
3536*	24	4.2	7.9	14257*	9	11.1	41.0
5437*	23	4.3	12.2	17689	8		
5774	22			18957	7		
7031	21			19765*	6	16.7	57.7
7996	20			19878	5		
8870	19			24708	4		
8892	18			25118	3		
9375	17			26490	2		
10114	16			29666	1		
10428	15						
11788	14						

* DENOTES FAILURE.

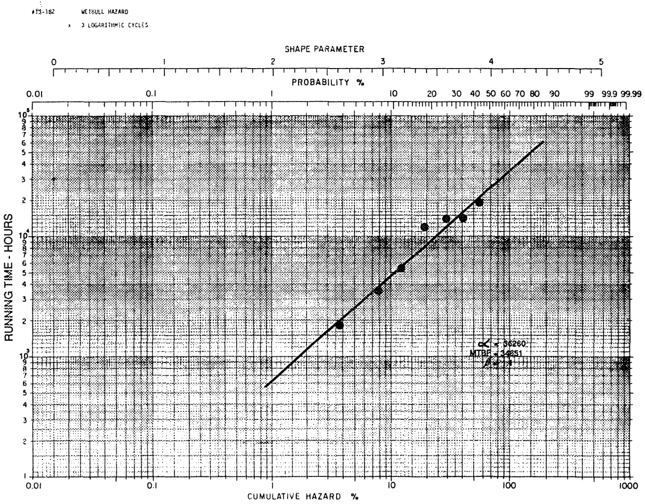

Figure 6. Case 1—Hazard Analysis of 21000 HP Motors with Running Time as the Parameter.

Because of this independence of failure on hr, an analysis was made based on the number of starts of the motors. Because of the large forces and heat generated by the inrush current on large motors, the total number of starts is frequently the most important factor in reducing the reliability of large motors. The failure date related to the number of motor starts (and of course shutdowns) is given in Table 3. This data is plotted in Figure 7. It can be seen that with a beta of 1.84, the failure mode is one of wearout where the failure rate increases with number of starts. The average number of starts before failure is 100. Comparing the differences in correlation of the two analyses, there seemed to be a strong indication that the problem was probably associated with starting (or stopping) the units, rather than with running or installed time.

Table 3. Hazard Analysis of Insulation Failure Data of 21000 HP Motor for Number of Starts.

STARTS	REVERSE RANK N	HAZARD 100/N	CUMULATIVE HAZARD	STARTS	REVERSE RANK N	HAZARD 100/N	CUMULATIVE HAZARD
14	28			51	13		
18*	27	3.7	3.7	57*	12	8.3	22.9
23	26			58	11		
24	25			59	10		
25	24			60	9		
25*	23	4.3	8.0	63	8		
26	22			64*	7	14.3	37.2
27	21			66	6		
31	20			75*	5	20.0	57.2
35	19			81	4		
38	18			105	3		
38	17			116*	2	50.0	107.2
41	16			148	1		
49	15	6.6	14.6				
50	14						

* DENOTES FAILURE.

With this insight, a motor was selected for testing with surge measuring instrumentation furnished by the motor manufacture. Of principle interest were measurements made during the

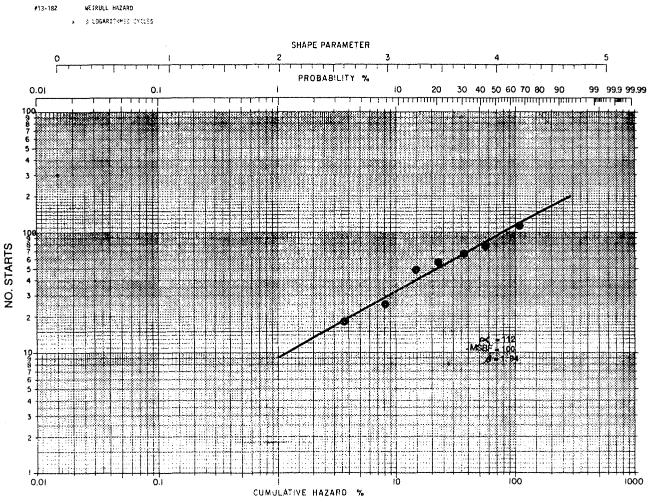

Figure 7. Case 1—Hazard Analysis of 21000 HP Motors with Number of Starts as the Parameter.

start and stop. It was discovered that during an aborted start (where the unit is tripped before reaching full speed) the air operated breaker restruck when opening and imposed a short duration, high magnitude voltage spike on the system. Although the voltage spikes were much higher than expected (indeed the breaker manufacturer did not think that it would restrike at all), the measured levels were not as high as the quoted breakdown strength of the insulation. It is suspected that some turn-to-turn insulation was weaker than the manufacturer expected.

To solve the problem, double insulation was installed on all rebuilt machines. This measure has been very successful with no failures on any of the rebuilt motors. However, the remaining original motors continued to fail and in Figure 8 a comparison is shown between the hazard analysis after the seven failures, and one done several years later with 20 failures. Of particular interest here is the accuracy of the estimated life with only seven failures (34651 hr) compared to that made later with 20 failures (29000 hr). The rather close agreement shown in the figure indicates that with a single failure mode, such as evident here, extrapolations can be valid even with a relatively small number of data points. This is very useful in studies where one does not have the time (or desire) to allow most of the units to fail.

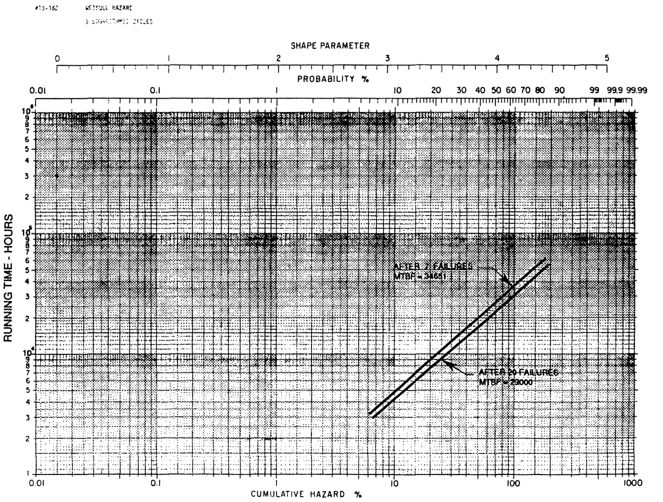

Figure 8. Case 1—Comparison of Life Estimate of 21000 HP Motors with Seven Failures and 20 Failures.

Case History 2—Mechanical Seal Failures

This case deals with a mechanical seal problem on a large number of multistage pumps in water injection service. The seals had several failure modes which could be identified by inspection. Among these were scale buildup due to the high salinity of the water, overpressure, and overheating. The data in Table 4 show how the failures were labeled for a hazard analysis. When the hazard data are plotted for all of the failures in Figure 9, several conclusions can be drawn. First the characteristic life is seen to be about 100 days. In this case, the MTBF is not used, since the data indicates by the sharp bend in the data that several failure modes are present, and one value for β is not applicable to the entire curve. It is seen that there appear to be three distinct portions of the curve which have β of 1.6, 0.5, and 1.2, respectively.

Table 4. Failure Data of Mechanical Seals with Multiple Failure Modes.

DAYS CODE	REVERSE RANK N	HAZARD 100/N	CUMULATIVE HAZARD	DAYS/ CODE	REVERSE RANK N	HAZARD 100/N	CUMULATIVE HAZARD
1 I	78	1.28	1.28	27 +	39		
1 I	77	1.3	1.58	27 P	38	2.63	54.7
1 +	76			31 +	37		
2 I	75	1.33	2.91	35 H	36	2.78	57.48
2 I	74	1.35	4.26	36 H	35	2.86	60.34
3 I	73	1.37	5.63	37 S	34	2.94	63.28
3 I	72	1.39	7.02	38 H	33	3.03	66.31
3 I	71	1.41	8.43	42 A	32	3.13	69.44
3 +	70			45 A	31	3.23	72.67
4 E	69	1.45	9.88	49 P	30	3.33	76.0
4 H	68	1.47	11.35	53 +	29		
4 U	67	1.49	12.84	63 A	28	3.57	79.57
4 U	66	1.52	14.36	68 D	27	3.7	83.27
5 A	65	1.54	15.9	70 P	26	3.85	87.12
5 E	64	1.56	17.46	81 P	25	4.0	91.12
5 U	63	1.59	19.05	74 P	24	4.17	95.29
5 E	62	1.61	20.65	84 P	23	4.35	99.64
5 P	61	1.64	22.29	84 +	22		
7 E	60	1.67	23.96	86 +	21		
7 H	59	1.69	25.65	90 P	20	5.0	
8 P	58	1.72	27.37	98 H	19	5.26	104.64
10 +	57			98 P	18	5.56	109.9
10 P	56	1.79	29.16	99 S	17	5.88	115.46
11 P	55	1.82	30.98	106 H	16	6.25	121.34
11 H	54	1.85	32.83	107 P	15	6.67	127.59
13 P	53	1.89	34.72	107 A	14	7.14	134.26
13 +	52	1.92		110 S	13	7.69	141.4
16 H	51	1.96	36.68	114 P	12	8.33	149.09
17 D	50	2.0	38.68	118 H	11	9.09	157.42
17 D	49	2.04	40.72	118 P	10	10.0	166.51
18 +	48			120 H	9	11.1	176.51
18 +	47	2.13	42.85	144 P	8	12.5	187.61
18 +	46	2.17	45.02	155 H	7	14.29	200.11
20 +	45			162 S	6	16.7	214.39
22 +	44			165 P	5	20	231.09
23 H	43	2.23	47.25	169 A	4	25	251.09
24 S	42	2.38	49.63	182 +	3		276.09
25 P	41	2.44	52.07	190 +	2		
27 +	40			202 +	1		

+ - NONE FAILURE
I - INSTALLATION
A - ABRAISION
P - OVER PRESSURE
H - HEAT
S - SCALE
E - EQUIPMENT
D - DEPOSIT

The first portion of the curve is believed to be an anomaly due to a systematic error in recording the running time. Because of the extreme isolated location of the injection pumps (they are located about 50 miles out in the desert), it is suspected that several days lapse can occur between the time of actual failure and the issuing of the work order to repair the seal. For this reason, the first portion of the data, β = 1.6, is distroted and results in an apparent (and incorrect) wearout mode. If a day or so is subtracted from the data, it is seen that the failures for the first several days will tend to fall on the line with a β of 0.5, a wearin mode. Thus, there seem to be a predominant wear in failure mode up until about 100 days of running, at which time the mode changes over to a wearout mode.

This case points out one of the powerful uses of a hazard analysis; predicting the effects of a design change. In this exam-

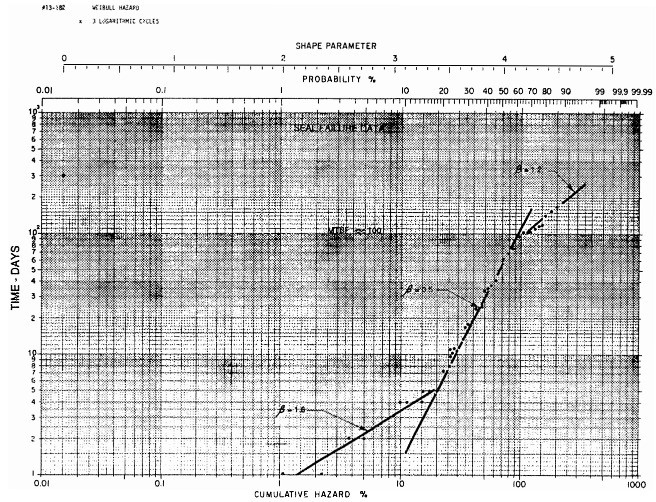

Figure 9. Case 2—Hazard Analysis of Mechanical Seals on Water Injection Pumps Showing Wear-In and Wear-Out Failure Modes.

ple, it was determined by engineering that one of the failure modes of the seal was overpressurization, caused by a leaking check valve when the pump was shutin. A potential fix was recommended of installing a pressure relief system at a cost of $40,000 per pump. The effects on realiability for this modification can be estimated with the hazard analysis by considering that the overpressurization mode was censored data. That is, the analysis was run with the overpressure mode considered as a nonfailure. When this was done, it was shown that even if the fix were 100 percent effective, the overall life of the seal would only be increased by only about 20 days. A simple economic analysis showed that this minor improvement in life could not justify the expenditure and the project was canceled.

Case History 3—Turbine Shroud Problem

This case describes the analysis of several shroud/pinion failures on the first stage of a 30,000 HP steam turbine. Of a population of eight turbines, there had been three failure of the tenons on the first stage shroud. Although only three points are not a very satisfactory sample size from a statistical point of view (the confidence limits tend to be rather wide), they are still enough to draw a line and can yield insight into the failure mechanism. The data for this case are shown in Figure 9 and indicates that with a β of 2.8, the failure mode is one of wearout and the MTBF is 19000 hr. A wearout mode is consistent with the fatigue failure that was identified on the tenons. However, this interpretation presents a problem in that, if the mode were one with an increasing failure rate, other machines in the population with up 43,661 hr should have failed also. Engineers are thus faced with a scenario in which there is a wearout failure of young machines that is somehow terminated. If a machine is able to reach a certain age (about 15,000 hr), its failures stop and it can reach a long and happy life.

Based on this analysis, it was hypothesized that such a mechanism did exist. Since damping in the wheel controls the amplitude and thus the stress, it was suggested that initially, in young machines, the blade shrouds and tenon were all penned tight with little relative movement and thus, low coulomb friction damping. They would, thus, be subjected to large potential stresses and fatigue. If, for some reason, the wheel did not fail and had a chance to wear and loosen up, the damping would increase and the stresses would be reduced.

Insitu measurement of the wheel damping were made on all of the turbines and plotted as a function of running hr [7]. This plot did in fact show that initially the damping in the wheel was low at about 0.5 percent of critical and after 40,000 hr had increased by a factor of five, to 2.5 percent. The manufacturer redesigned the shroud system which increased the initial damping by a factor of about three. The new design is currently running satisfactorily and there have been no more failure of the old machines.

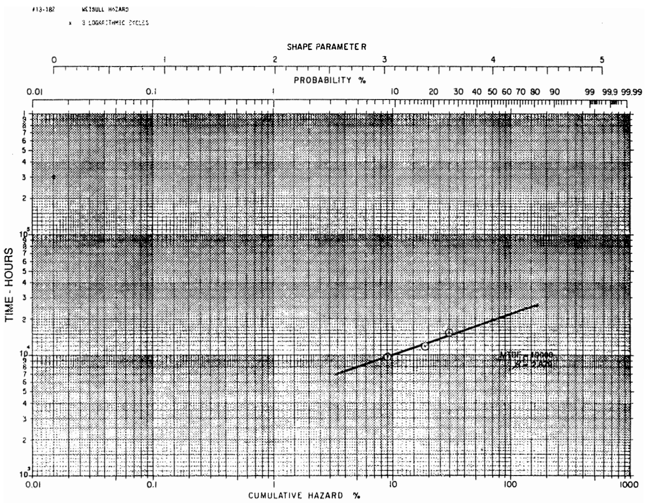

Figure 10. Case 3—Hazard Analysis of Turbine Shroud Failures Showing Wear-Out Failure Modes.

Case History 4—Crude Shipping System Motor Problem

This case describes the analysis of several failures on the motor drivers of a crude oil shipping system. The population consisted of eight induction and synchronous motors of 6000 to 8000 hp. The motors were in a very severe service since the shipping system required that they be started and stopped several times a day. Some of the machines had accumulated over 2500 across-the-line starts. The data for five first time failures are plotted in Figure 11. This hazard plot shows that there is a wearout mode with a βof 5.6. This represents a rapidly increasing failure rate with a charactistic life, α, of about 40 months.

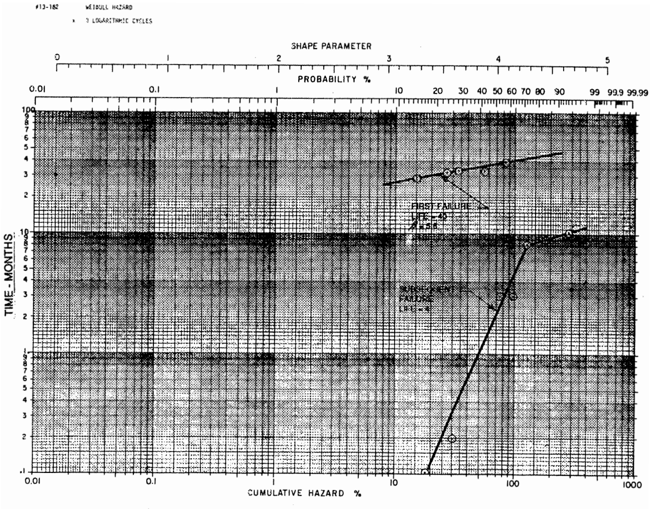

Figure 11. Case 4—Hazard Analysis of Shipping System Motors.

In this population, three of the machines had never failed, although their installed times were between 36 to 45 months and, thus, were well into the wearout range. Using the technique described above to calculate probability to failure, it was estimated that there was a 95 percent probability for one of the three machines failing in the following five month time period. It was, therefore, recommended that the three remaining motors be opened for inspection since they seemed to be in a very vulnerable region of their life cycle. Considering the hazard technique to be highly suspect, Operations very reluctantly agreed to inspect only one of the machines. This machine was found to have a bolt backed out that held the exciter strap to the shaft. This fault would have caused a failure of the strap in a very short period of time and resulted in a trip to the shop. Instead the bolt was reset and locked in place and the motor returned to service. With no further discussion, the remaining machines were then inspected and found to also have problems.

Also shown in Figure 11 is the analysis of succeeding failures that occurred after the initial ones. It is seen that with these subsequent failures there is a drastic decrease in life, from 40 to four months, and that there is a distinct infant mortality period. These data indicate that the repairs were not bringing the motors back to a "like new" condition and that there seem to be a problem with installation. These concerns were reviewed with the maintenance organization and the problems were addressed.

The root problem of frequent starting and stopping was solved by redesigning the system so that the motors, which drove through a variable speed coupling, could remain running even though the pumps were stopped.

Case History 5—Water Injection Pump Impellor Cracking

This case concerns a piping modification whose primary purpose was to correct a subsynchronous vibration problem on a pump [8]. The 20,000 hp pump was in water injection service and, in addition to the vibration problem, also had a history of impellor vane cracking.

The installation of long taper piping transitions on the suction and discharge of the pump cured the vibration problem and seemed to yield some benefit to the vane cracking problem. However, examination of the impellor casting showed that there were major variations in vane thickness and the manufacturer was requested to replace them with precision cast impellers with more uniform, thicker vanes. The manufacturer's response was that the piping change had cured the cracking problem and there was no need for a change in impellor design and manufacture.

A hazard analysis was performed on the impeller failure data for the period both before the piping change and after the change. This analysis is shown in Figure 12. As seen here, the MTBF before the change was 2500 hr, and after the change was 6000 hr, an improvement of 240 percent. There had indeed been a significant gain. However, as was pointed out to the manufacturer, 6000 hr was still not good enough. The impellors were replaced with no further discussion needed.

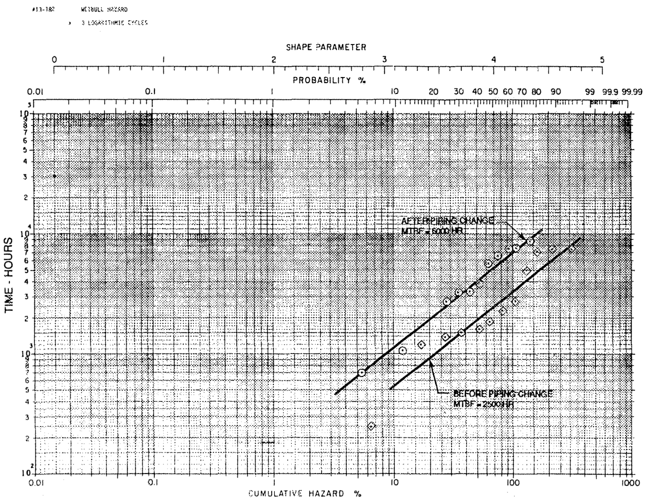

Figure 12. Case 5—Hazard Analysis of a Pump Impellor Failures Showing Differences in Life Before and After a Piping Modification.

Case History 6—Gear Coupling Life

This case was a simple one of determining the life and failure mode of a population of 12 large (6000 to 16000 hp) gear type couplings on motor drives. As part of the effort to improve the reliability of the crude oil shipping system described in Case 4, it was found that there had been three failures on the gear couplings, and these were a factor that reduced the system operating factor.

A hazard analysis was performed for the coupling data, and the results are shown in Figure 13. It is seen that with a beta of 2.97 the couplings were in a wearout mode and had an average life of 42 months which was not acceptable. Examination of a failed coupling showed that the coupling grease was dry and hard, and records indicated that the grease had probably not been changed since startup. This type of failure would produce a wearout mode which was in agreement with the conclusions reached from the hazard analysis. In addition, the grease being used was not one recommended by the coupling manufacturer as being acceptable for this service. To correct the problem, a yearly regreasing schedule was set up and the company's lubrication specification was modified to include a coupling grease.

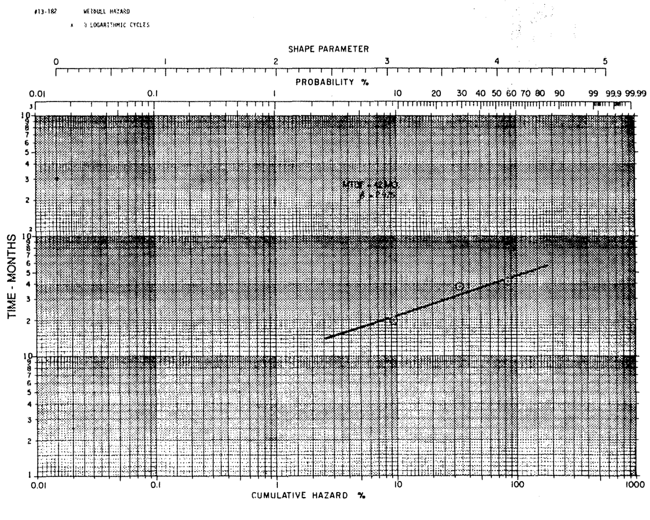

Figure 13. Case 6 —Hazard Analysis of Gear Type Couplings.

Case History 7—Electric Motor Rotor Failure

This case involves the analysis of failures of a population of five, 15000 hp, four pole induction motors. Within a year of startup of a large sea water treatment plant, failures began on the rotors of motors driving the main water shipping pumps.

After four failures, a hazard analysis was performed on the data. Like Case 1, the data indicated that the failures did not correlate well with time in service. The running hr plot indicated that there are two modes, even though only one mode existed, a burnout of the rotor bars. However, when the data was analyzed based on the number of machine starts, there was a very good correlation, with all of the data falling on a straight line. This can be seen in Figure 14, where the curve fit line has a β of about 2.0, indicating a wearout mode and the MTBS for the units was 80 starts. Even though there are only four data points, the confidence band is acceptably narrow due to the good correlation with number of starts.

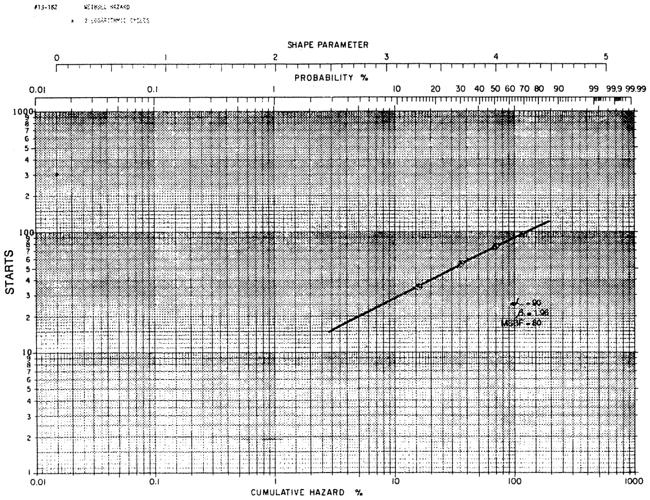

Figure 14. Case 7—Hazard Analysis of a 15000 HP Induction Motors with Number of Starts as the Parameter.

Based on the above analysis, it was concluded that the failures were caused by starts and not the running hr. The manufacturer was notified of these findings and conclusion. A thermal analysis by the manufacturer of the rotor bar design indicated that the problem was due to thermal fatigue caused by starting. The fatigue analysis estimated that the bar would fail in 100 starts which was in excellent agreement with the findings of the hazard analysis. The manufacturer recognized the design fault and supplied 10 new rotors of a modified design. These have been running successfully for nine years.

Because the time required to ship the motors back for repair, which would take about 10 months, there was a concern about the probability of having additional failures in that time. An option would be to retain motors for a temporary repair onsite. Based on an analysis of the probability of failures, which indicated only a small chance of there being enough failure to affect production in the 10 month time span, it was decided by management to ship all of the damaged machines back to the factory.

Case History 8—Mechanical Seal Improvement Study For Refinery

The last case involves a program to improve the life of mechanical seals in pump service for a refinery. This study encompasses a representative sample of 173 pumps in the refinery with a population of 244 seals. The major goal of the study is to demonstrate a five year average life of the seals on the program. To keep track of the seals, two data bases are maintained. The first has design data about the pump, process information, and the seal design information. The second data base records seal failure information such as running time, workorder costs, and description of the failures. All of this information can then be linked as required with the relational feature of the data base.

The use of the hazard analysis is crucial to the success of the program in that it allows the MTBF to be calculated not only for the plant as a whole but as a function of any parameter in the data base. Thus, the life of a seal as a function of location, plant, process, material or manufacturer can easily be calculated. When the program highlights a problem, and the seal design is modified to correct the problem, the effects of the change can be ascertained through the use of the hazard analysis.

The Weibull plot for all of the seals on the program is shown in Figure 15. This plot, which is updated monthly, makes it possible to quantify the progress of the program and report to management on a regular basis. As the plot shows, the MTBF for the refinery is currently 3.45 years, which is up from 2.9 years when the program started a year ago. With a β of 0.88, the overall data indicates there is a slight wear in or infant mortality. This data implies that an effort should be made to improve the seal installation procedures.

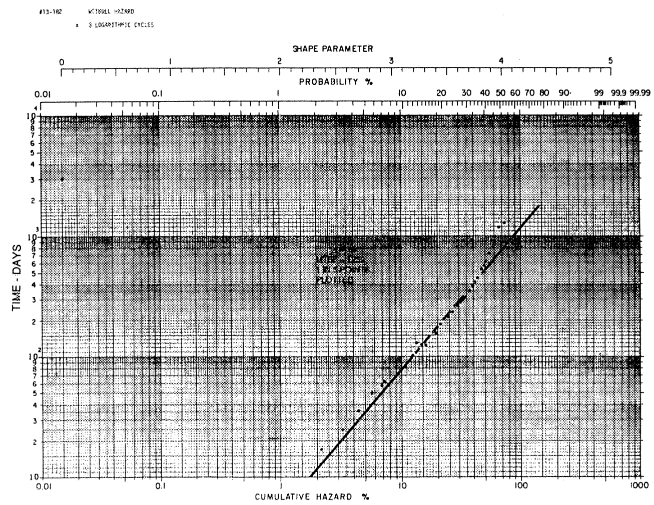

Figure 15. Case 8—Hazard Analysis of Refinery Pump Seals for 173 Pump and 244 Mechanical Seals.

One of the parameters affecting the life of seals that was analyzed, was the differences in life between silicon carbide and tungsten carbide hard faces. This analysis is shown in Figure 16, and indicates that the silicon carbide is significantly better than the tungsten carbide. Care should be taken, however, in interpreting the data since there is a significant extrapolation re-

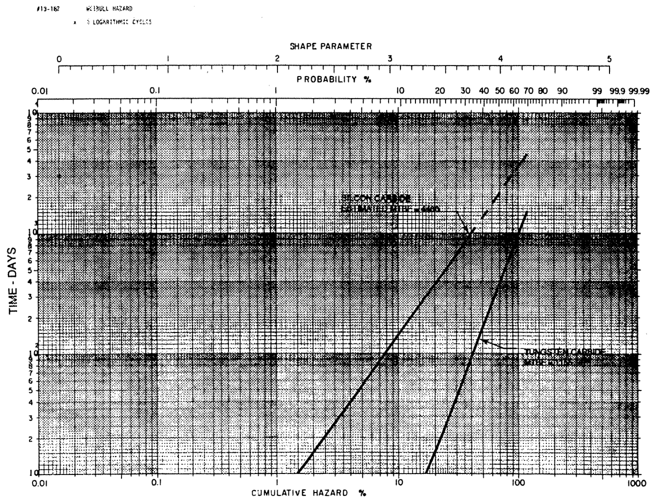

Figure 16. Case 8—Analysis of Seal Hardface Material Showing Comparison Between Silicon Carbide and Tungsten Carbide.

quired with the silicon carbide data. Although the analysis suggests a four to one improvement, there could be a wearout mode beyond the range of the data which could invalidate part of the extrapolation.

OTHER USES OF THE HAZARD ANALYSIS

In addition to the failure analysis shown here, there are many other uses that can be made of the technique that are not covered in detail here. By treating the data as simply an array of numbers, and not failures, a probability distribution can be obtained for any set of numbers. For example, the technique has been used to define the amplitude probability distribution of a random vibration signal. The peak values of a sample random time trace were input to the analysis. The resulting probability distribution from the Weibull plot of the data was then used successfully in a Monte Carlo routine.

In another example, the deviation of pump performance from the manufacturer's published curve was analyzed as a hazard problem. The analysis showed that while four of the pumps indicated that the performance followed a normal bell curve ($\beta = 3$), the point for one pump did not fit on the line. This pump was opened up and found to have an incorrect diffuser.

CONCLUSIONS

The hazard analysis has been found to be an extremely powerful tool in analyzing failure data of the type that is common to the turbomachinery industry. With a simple graphical procedure, it offers valuable, and in some cases crucial, insight into failure mechanisms and a statistical correct method to quantify the life of machines.

ACKNOWLEDGMENTS

The author wishes to thank Gary Darby of Intellect Systems, Incorporated and Saudi Aramco for his writing the hazard analysis computer program.

REFERENCES

1. Nelson, W., *"Applied Life Data Analysis,"*New York: John Wiley & Sons, Inc., (1982).
2. Nelson, W., "A Method for Statistical Hazard Plotting of Incomplete Failure Data that are Arbitrarily Censored," General Electric Research and Development Center, Schenectady, New York, TIS Report 68-C-007 (January 1968).
3. Bianco, E. G., "Reliametrics Notes," unpublished, General Electric Company, Appliance Park, Louisville, Kentucky.
4. Nelson, W., "Hazard Plotting for Incomplete Failure Data," Journal of Quality Technology, *1,* (1) (January 1969).
5. Nelson, W., "Hazard Plotting Methods for Analysis of Life Data with Different Failure Modes," Journal of Quality Technology, 2, (3) (July 1970).
6. Team - Special Purpose Graph Papers, Box 25, Teamworth, New Hampshire.
7. Corley, J. E., Campbell, W. R., and Graham, M. K., "In-Situ Measurement of the Damping in a Steam Turbine Wheel," Proceeding of the Vibration Institute (1988).
8. Corley, J. E. "Sub-synchronous Vibration in a Large Water Flood Pump," *Proceedings of the Seventh Turbomachinery Symposium*, Turbomachinery Laboratory, Department of Mechanical Engineering, Texas A&M University, College Station, Texas, p. 103 (1978).