A RELIABILITY ANALYSIS OF PERSONNEL PROTECTION SYSTEMS AT THE SPALLATION NEUTRON SOURCE

ADAM SPANNAUS

Department of Mathematics, University of Tennessee Knoxville, TN
and
Oak Ridge National Laboratory, Oak Ridge TN[1]

ABSTRACT: In this paper we present and analyze field gathered system reliability data from the Oxygen monitoring and Access Control systems managed by the Protection Systems team at the Spallation Neutron Source at the Oak Ridge National Laboratory. The time span of the data varies from roughly fourteen years for systems in place as the linear accelerator was being built and tested, to a few months for the newest instrument beam lines.

We analyzed the time to fail for the above systems, accounting for censored data, and developed a non-parametric probability model of the overall reliability for each system. From this probabilistic model, we were able to estimate the survivor and hazard functions, and find good correlation between these estimates and their empirical counterparts. Moreover, we used our model to find whether or not each system's life distribution is *New Better than Used*.

1. INTRODUCTION

Particle accelerator safety system design and management are heavily reliant on reliability and assurance principles found in other high-consequence technologies such as aerospace, chemical processing and nuclear industry [1]. High-demand, high-reliability systems present a unique set of challenges in both design and maintenance. As the Spallation Neutron Source (SNS) grows and seeks to add a Second Target Station, the systems will increase in complexity in addition to the demands on the system. A systemic failure can lead to not only loss of experiment time, but the potential for injury to personnel and damage to equipment.

The SNS is a linear accelerator at Oak Ridge National Laboratory that creates high-intensity neutrons for study by visiting and laboratory scientists. The Protection Systems monitor the safety and integrity of various safety systems throughout the facility primarily designed to keep personnel safe. Two of these systems are the focus of this study. Both systems were constructed as the linear accelerator and facility grew, beginning in the early 2000s. We have accumulated over 1.8 million hours on the programmable logic controllers and about 2.8 million hours on the Oxygen monitoring system.

The Oxygen Deficiency Hazard monitoring system (ODH) is a series of sensors that detect the level of oxygen present in the atmosphere. If the concentration of oxygen drops below 19.5%, the system sends an alarm signal indicating an oxygen deficiency. The oxygen level is of primary importance to personnel working in or around the facility, but can also signify problems with other systems.

The ODH system has oxygen monitors in place throughout the linear accelerator to the instrument beam lines where neutrons are detected. The system is designed with some redundancy, so that if one sensor fails the entire system is not taken off-line. Consequently, we have a large dataset of failure times and modes. The oxygen monitoring sensors are replaced yearly, so the corresponding dataset is right-censored. The oxygen monitors are subjected to a variety of conditions throughout, such as high radiation and low oxygen, that cause premature sensor failure.

The second system studied is the Access Control and Interlock system. All Access/Interlock input-output signals are routed through Programmable Logic Controllers, PLC, which allow or disallow actions based on the overall system's state. These controllers are installed redundantly, so if one PLC fails, the entire system remains operable in a safe state. PLCs fail infrequently, consequently, we have a small amount of data about their failure times and modes. It is important to note here the distinction made between the PLC itself and the software running on the PLC.

Through careful analysis of the failure data, we hope to become increasingly proactive in preventing costly system failures. Currently these systems are designed with the manufacturer's mean-time-to-fail, MTTF, in modeling for a failure analysis. To increase the system reliability, understanding the failure rate and modes on site is imperative. For the scope of this paper, we limited our study to the ODH and Access Interlock systems because of the different set of challenges associated with each.

When considering failure modes and time to fail with regard to the Access Interlock system, we must be careful to separate the software from the hardware components. It is not software that fails, but the design requirements [2]. In February 2011, the SNS had a Protection Systems event that was significant enough to shut off beam production. The safety system worked as it was designed, but the root cause was discovered to be faulty firmware on the instrument PLCs. The problem was that in high-demand, continuous applications, such as the SNS, the PLC would go to a safe state after three to six months of continuous operation and could not be restarted. Failures related to software performance are not a failure of the device itself, but are of a different type. Such failures are the subject of much research, and are not considered here. As a direct result, the PLC time-to-fail data was not independent, and we censored some data points to have a more accurate picture of the component's time-to-fail, instead of the software. Hence there exists a strong correlation between subsequent changes to the PLC system and this specific event.

From developing an estimate of each system's life distribution, we seek to find if the distribution is New Better than Used (NBU), that is akin to the idea of 'Positive Aging.' So, if a component has a lifetime, t_i, and if the system fails when a component fails, then the overall system's lifetime is T_i, where i denotes the lifetimes for each individual component [3]. As a practical application, the Protection Systems team (PST) will use this information in the planning and design of new systems.

In §2 we briefly describe our analysis of the failure data for both systems, and develop some necessary concepts. In §3 we present the findings of our failure analysis, and whether the distributions are NBU. We follow with a brief summary and considerations for further investigation in §4. In all cases following, X is a random variable, $f(x)$ represents a life distribution, $F(x)$ its density function, $\bar{F}(x) = 1 - F$, and $\lambda(x)$ is the hazard function.

2. METHODOLOGY

Our investigation began with the hypothesis H_0 that each system's lifetime is drawn from a Weibull distribution. We chose this distribution since it frequently provides a good model for reliability analysis [4]. We tested the hypothesis H_0 vs. H_1, that the lifetimes are not from a Weibull distribution. Upon visual examination, we found the distribution to be multi-modal. Additionally, upon further examination, we found four distinct failure modes for the ODH system. The likelihood of all four failure modes sharing the same parameters is low [5], so we rejected H_0, and instead took a non-parametric approach to the data analysis.

Non-parametric methods do not have any fixed structure aside from the kernel used to estimate the density. Instead, they rely upon the data to determine an estimate and smooth the contribution of each data point over the neighborhood of the point [6]. The contribution of each data point x_i at a point $\hat{x}$ is dependent on how close x_i and $\hat{x}$ are [7]. The kernel determines the contribution and shape of the kernel function at each point. The distribution function is estimated by the following:

$$\hat{f}_h(x) = \frac{1}{nh} \sum K(\frac{x - x_i}{h})$$

where h is the bandwidth. We used the Epanechnikov kernel, since it is optimal as compared to other kernel density estimators in the Mean-Squared error sense [6] [8] . It is given by: $K(u) = \frac{3}{4}(1-u^2)\mathbf{1}_{\{|u|\leq 1\}}$, where $\mathbf{1}$ is the indicator function [6]. In Kernel density estimation, the most important determination is the kernel bandwidth, which determines the amount of smoothing for the density estimator. If the value is too high, the density estimate is over-smoothed and obscures some of the data's underlying structure. Similarly, choosing a value that is too low will lead to undersmoothing and a density estimate with many spurious points [6].

The implementation of the kernel density estimations were programmed using Matlab. We compared two kernel density estimations: our own function and the built-in `ksdensity` method in Matlab. The differences between the two are not subtle, even with the same bandwidth. The `ksdensity` function assumes the data to be from normal distributions [9]. This produces a smoother density estimate, but at the cost of obscuring some of the data's shape. Our implementation however, while not as smooth, gives a better picture of the life distribution, since we do not making any assumptions about the distributions of the data. We do however use the bandwidth estimate given by Matlab. Typically in kernel density estimation, the bandwidth parameter is chosen to minimize the mean squared error between density estimations [7]. In a further investigation, optimizing this parameter to the data would produce a better result, one tailored to this data. We then accept or reject the density estimate based on a visual inspection, verifying the estimate matches a histogram of the data.

From an estimate of the probability density, we then constructed an estimate of the survival function [10] and compared it against the empirical survival function. We determined the empirical survival function from the data using the Kaplan-Meier estimator, which is defined as follows:

$$\bar{S}(t) = \prod_{t_i < t} \frac{n_i - f_i}{n_i}$$

where n_i is the number of the population at risk in each time interval, t_i, and f_i is the number of failures in each time period [10] [11]. $\hat{S}(t)$ gives a point estimate of the survival function, accounting for censoring of the data, specifically right-censored data [10]. Right-censoring of data is when a member of a population is removed from a study with an indefinite end result. Accounting for censoring is important when finding a failure rate; since the censored population has not failed, a precise value cannot be known.

Once we determined an estimate for the survival function, we then found the empirical cumulative distribution and compared it with the cumulative kernel density estimate. To do so, we constructed point

estimators for the survival function using the Kaplan-Meier estimator, and its complement, the cumulative density function [10]. As shown below, our estimate of these functions falls within the 90% confidence interval for the empirical function. We also provide the same estimates with 60% confidence bounds shown, which is typical of reliability studies [4]. While we do not have an exact correlation at the 60% confidence level, the fit is still quite good.

Once the density and life distributions were determined, we then found the hazard and cumulative hazard functions. The cumulative hazard function is given by

$$-\ln[1-F] = \int_0^t \lambda(u)\ \mathrm{d}u$$

and by the Fundamental Theorem of calculus, we may uniquely determines the probability mass function [12], as shown below.

Let F be a density function, with life distribution f, survival function $\bar{F}$, and hazard function $\lambda(t)$. Then letting $G'(t) = -\ln[\bar{F}(t)]$, we find:

$$\begin{aligned} G'(t) &= \frac{f}{\bar{F}} = \lambda(t), \\ \int_0^t G'(u)\ du &= \int_0^t \lambda(u)\ \mathrm{d}u, \\ G(t) - G(0) &= \int_0^t \lambda(u)\ \mathrm{d}u, \\ -\ln[\bar{F}(t)] &= \int_0^t \lambda(u)\ \mathrm{d}u. \end{aligned} \tag{2.1}$$

It is this characterization of the cumulative hazard function that allows us to determine if F is New Better than Used (NBU). To show this, we first need to prove the following proposition.

Proposition: If a function f, $f \in C[a,b]$, $f : \mathbb{R}^+ \longrightarrow [a,b]$, is convex and increasing on $[a,b]$, then f is superadditive on $[a,b]$.

Proof. By definition, $f(\alpha x + (1-\alpha)y) \leq \alpha f(x) + (1-\alpha)f(y)$ convex on $[0,1]$ $\forall \alpha \in [0,1]$ [13].
It follows then $\forall x, y \in [a,b]$,

$$\begin{aligned} f(x) + f(y) &= f((x+y)\frac{x}{x+y}) + f((x+y)\frac{y}{x+y}) \\ &\leq \frac{x}{x+y}\ f(x+y) + \frac{y}{x+y}\ f(x+y) \\ &= f(x+y). \end{aligned} \tag{2.2}$$

which is the definition of superadditivity. □

Thus it suffices to show if $\lambda(t)$ is increasing in $t, 0 \leq t$, then F is NBU. To show this, we consider the function $f = -\ln[1-F]$ in the above inequality. That is, if $f = -\ln[1-F]$ is convex, then $-\ln[1-F]$ is superadditive, and moreover, F is an increasing failure rate (IFR) distribution. This is equivalent to the hazard function, $\lambda(t) = \frac{f}{\bar{F}}$ being increasing in t. This follows directly, since a necessary and sufficient condition for a function to be convex on an interval $[a,b]$ is that $f(x) \geq 0, \forall x \in [a,b]$. Now since $\ln(x) \in C^\infty$, we find that

$$\frac{\mathrm{d}}{\mathrm{d}t}\ln[\bar{F}] = \frac{\mathrm{d}}{\mathrm{d}t}\int_0^t \lambda(u)\ \mathrm{d}u \tag{2.3}$$

$$= \lambda(t) \tag{2.4}$$

It remains to show that $-\ln[1-F]$ is superadditive. This condition can be satisfied by showing $\lambda(t)$ is monotonically increasing in t. However, since we do not have a closed form of this function, we instead consider divided differences to numerically approximate the derivative. That is where the value of

$$\frac{\mathrm{d}}{\mathrm{d}t}\lambda(t) \approx \frac{\lambda(t_i) - \lambda(t_{i+1})}{t_i - t_{i+1}} \tag{2.5}$$

is strictly greater than 0. Where $\frac{\mathrm{d}}{\mathrm{d}t}\lambda(t) > 0$ are the subsets where the cumulative hazard function is increasing in time. Then F is NBU, since a distribution having an increasing failure rate (IFR) is a sufficient condition for the life distribution being NBU.

3. RESULTS AND FINDINGS

The SNS has various document tracking applications simplifying the task of searching for events. From the observed and documented failures, we determined the mean time between failure for each system. The SNS uses multiple data repositories: Data Stream, a database of all work orders for work done at the SNS, and ProjectWise, project collaboration software. The SNS also keeps an electronic logbook, in which all events are recorded and labeled by priority level and according to which groups, areas, and systems are affected. Lastly, the PST group keeps a historical archive of equipment failures and maintenance records for their systems. Once collected and organized, the data was cross-referenced for causality and duplication.

The most common assumption for failures of electronic components is they have a Weibull distribution [5], but as shown in the figures below, this model is not a good fit for the systems studied here. A visual analysis of a Weibull plot of the failure times confirms the failure-rate distributions are multi-modal. As seen in figures 1 and 2, if the failure rates were from a unimodal Weibull distribution, the times would be almost linear [4]. Instead, the failure rate for the PLC system is a type E curve, and the ODH is type D, as classified by Jiang and Kececioglu [14].

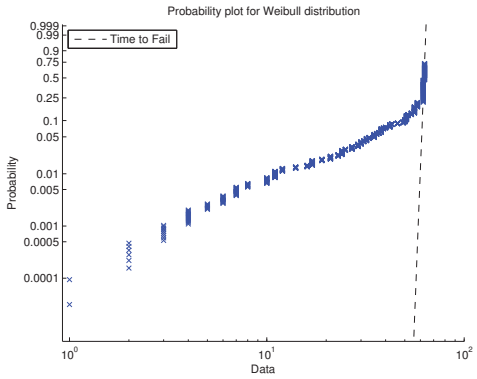

Figure 1: Weibull plot of ODH time to fail

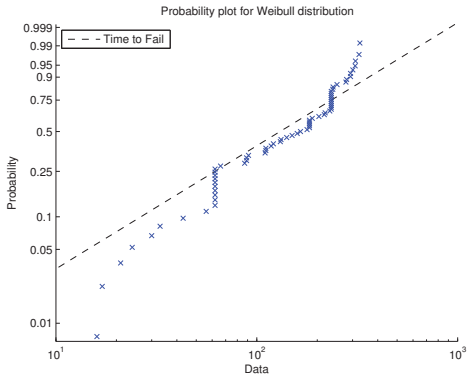

Figure 2: Weibull plot of PLC time to fail

From these data points we constructed point estimators for the survival function using the Kaplan-Meier estimator, and its complement, the cumulative density function. As shown in figures 3 and 4, our estimate of the survivor function falls within the 90% confidence interval for the empirical function.

Figures 5 and 6 are the cumulative density estimates with 90% confidence bounds shown. While we do not have an exact correlation, the fit is still sufficient for this application.

Both of the life distributions can be interpreted in relation to the Bathtub curve used to evaluate overall reliability [4]. The bathtub curve has an initially high failure rate during a break-in period, then the rate

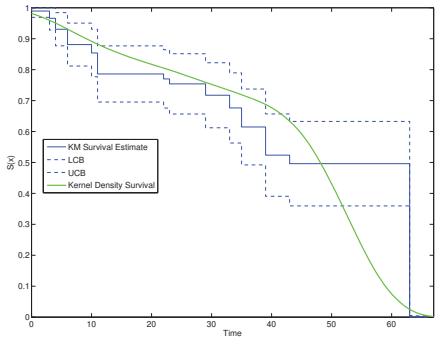

Figure 3: ODH Survivor function and density estimate

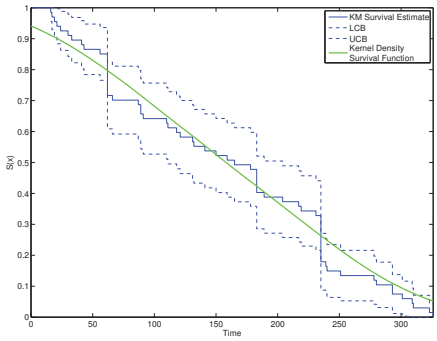

Figure 4: PLC Survivor function and density estimate

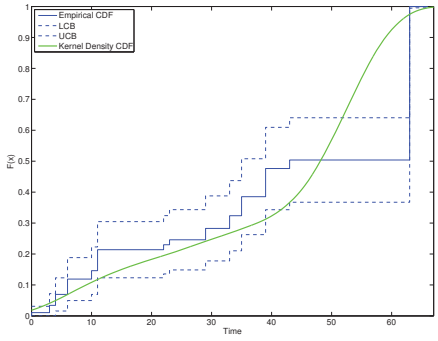

Figure 5: ODH empirical cdf and density estimate

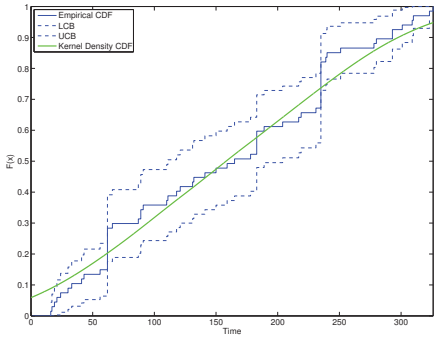

Figure 6: PLC empirical cdf and density estimate

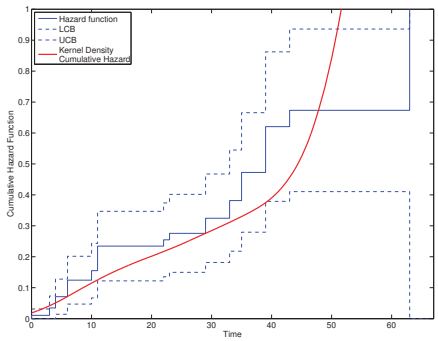

Figure 7: ODH empirical cumulative hazard function and Hazard function estimate

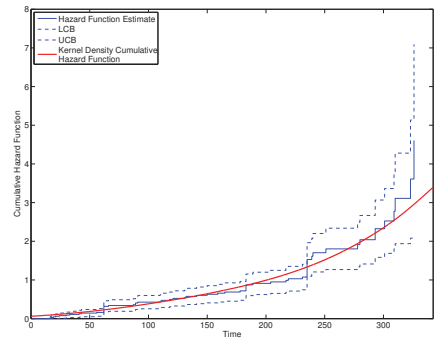

Figure 8: PLC empirical cumulative hazard function and Hazard function estimate

decreases to an almost constant failure rate. Lastly there is a wear-out period when components begin to fail, and the failure rate accordingly increases. How this relates to the problem posed here is that if an ODH survives to week 6, the failure rate decreases until we replace all the sensors around the 52 week mark. This event is especially evident on the survivor graph where the trendline goes rapidly towards zero. This correlates with the time when all of the ODH sensors are replaced, whether or not they have failed.

The causal effects on the PLC plot is less clear though. The rapid increase towards the 50 week make corresponds to the first certification of a system after commissioning. It is at this point that a majority of the errors in the system are caught and fixed, thus causing the increase in the hazard rate. The rate then decreases and remains roughly constant until something is changed in the system. Some examples of changes we found are: processor upgrade, change in the logic/program, or a firmware update.

Recall that the kernel density estimation is very sensitive to our choice of bandwidth [6]. If the bandwidth is too small, the result could be interpreted as a constant failure rate, which simply by visual examination, is obviously false. We chose to use the built-in Matlab bandwidth estimator, which is 'good enough' for most [4] applications. It is certainly a topic worthy of further investigation to determine an optimal bandwidth for the density estimates presented here. Following in figures 9 and 10 are the kernel density estimates with the histograms of the times to fail for each system.

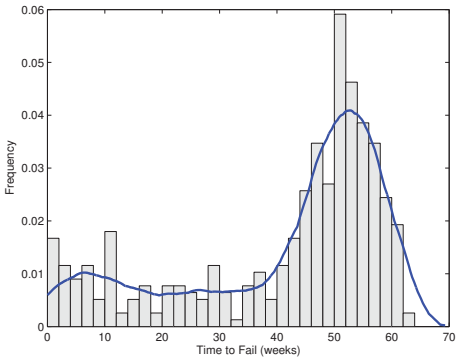

Figure 9: ODH density estimate

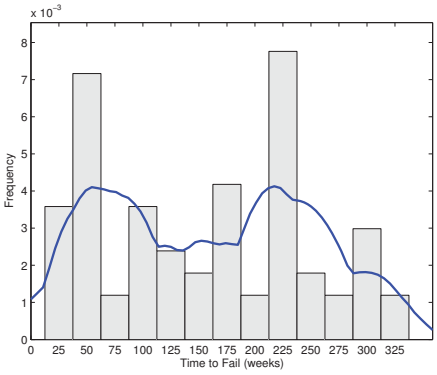

Figure 10: PLC density estimate

Once the density and hazard functions have been estimated, we seek to find if the systems are NBU or New Worse than Used. We evaluate $R(t_0, t_1) = \frac{f(t_0)-f(t_1)}{t_0-t_1}$, where $f = -\ln[1-F]$. Recall that $-\ln[1-F] = \int_0^t \lambda(u)\,\mathrm{d}u$, which is the cumulative hazard function. Hence, where $R(t_0, t_1) < 0$ we have that F has a decreasing failure rate, and similarly when $R(t_0, t_1) > 0$, F has an increasing failure rate. We then seek the inflection points of $R(t_0, t_1)$, which are precisely the critical points of $\lambda(t)$.

In our Matlab function, we wrote a simple algorithm to find where the function $R(t_0, t_1)$ changes sign. What we find is the ODH system is NBU from weeks 7 to 34, and then again from week 66 to 68. This last interval is not particularly useful however, since the manufacturer's recommended lifetime for an oxygen sensors is one year. The PLC are NBU for 338 to 347 weeks. These intervals agree with our intuition upon consideration of the histograms representing each system's time-to-fail.

Once we developed the above test for the NBU property, we applied the test to individual components of the ODH system. From the collected data, we were able to give an estimate of the life distribution. Then examined two different subsets to determine if the life distribution was NBU: (1) consecutive time intervals where the hazard rate is increasing and (2) time intervals from $t = 0$. These subsets were chosen for their direct application to maintenance of ODH systems. Consecutive time intervals give insight as to how the component is aging, and if the likelihood of a failure is increasing. Secondly, comparing the current lifetime of a component against $t = 0$ compares the current hazard rate against that of a new component. The results of the test as applied to four oxygen monitoring sensors is given below.

Figure 11 shows a 'normalized' hazard plot where we normalized the hazard function of sensor 204 against the hazard function of sensors: 207, 208 and 209. This plot shows how the associated hazard for each sensor is changing in time, relative to sensor 204. We chose this sensor since it has the longest mean

ODH Sensor NBU Comparison

Sensor Number	Consecutive Weeks	Weeks from $t = 0$
204	1 - 68	1 - 68
207	1 - 64	64 - 68
208	1 - 44	1 - 53
209	1 - 48	1 - 63

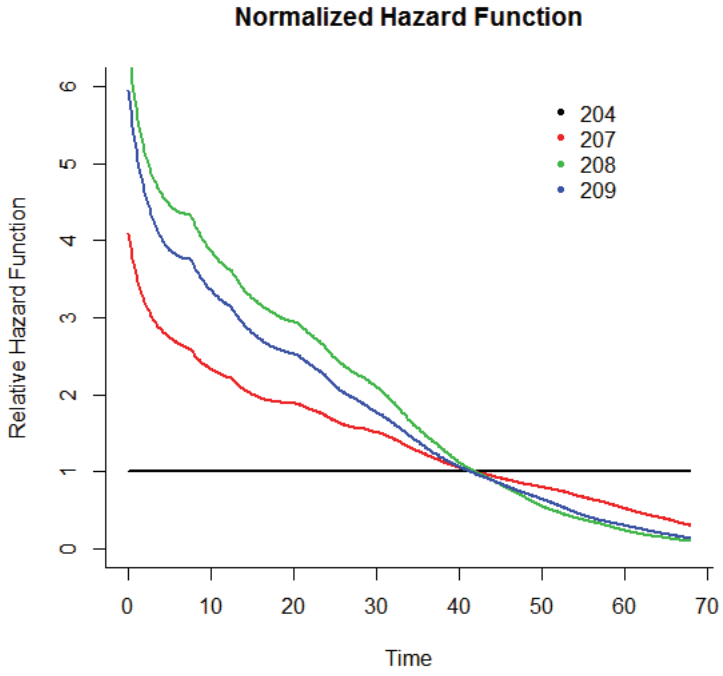

Figure 11: Normalized Hazard Function

lifetime of any sensor in the liner accelerator. The hazard rate for sensor 204 increases very slightly until $t \approx 40$, which intuitively makes sense, since these sensors are replaced on an approximately yearly basis. We also observe that the value of the hazard functions for the other sensors are less than that of sensor 204 after $t = 40$. This makes sense, since the few of sensors 207, 208 and 209 that had lifetimes greater than 40 weeks, were significantly longer. However, these long-lived sensors were in the first few years of the facility, before the linear accelerator was running at the current power levels.

4. CONCLUSIONS and FURTHER INVESTIGATIONS

We employed kernel density estimation to find the life distributions of each system, and from the density estimation, we were able to fit our estimates of the hazard and survivor functions against the empirical functions taken directly from the data. Having this information allows us to find the probability that a system will survive beyond a certain time. Such information is useful in Hazard and Operability Study (HAZOP) where specific hazards are identified, their likelihood assessed and the severity of each hazard [4]. At the SNS, HAZOP studies are used in the planning stages of new construction or as a tool to improve our safety and reliability.

We wrote a function to evaluate the Epanechnikov kernel, within a larger Matlab routine. This program dynamically allocates space as needed, and will update its values as more data points are added to the database. This database and Matlab routines will be used in the future planning for new installations, system modifications and in other reliability studies of the PST systems.

We developed a new test for NBU by using the shape of the cumulative hazard function and finding where it is convex, which implies superadditivity and inclusion in the NBU class.

From estimating the cumulative hazard and survivor functions, we were able to determine in which time intervals the system was NBU. Such information will allow the PPS team to better plan and make predictions about the systems they design and manage. These models will be especially useful in developing the systems for the Second Target Station at the SNS, which is in its development phase.

To improve the accuracy of this study, we should choose a more robust method to determine the kernel bandwidth, either via cross-validation, as proposed by Wasserman [7], or the method given by Raykar and Duraiswami [15]. As our density estimate is heavily reliant on the choice of bandwidth [7], these are a few promising options to improve the accuracy in further studies, specifically through a parameter estimation based on empirical data.

The method proposed by Raykar and Duraiswami, gives an ϵ-*exact* approximation to the univariate Gaussian kernel. Alternatively, Wasserman presents a method based on minimizing the cross-validation estimator of risk [7]. That is minimizing the function:

$$\hat{J}(x) = \int \hat{f}^2(x)dx - \frac{2}{n}\sum_i \hat{f}_{-i}(X_i)$$

where $\hat{J}$ is the estimator of the cross-validation score and $\hat{f}_{-i}$ is the kernel estimate omitting X_{-i}. By minimizing this function we can find an optimal choice for the bandwidth in terms of the mean-squared error[7].

Additionally, a similar test could be done for individual components of each system, thus strengthening our assertion of the NBU intervals [16]. This test is described in a paper Proschan, one of the authors of the New Better than Used paper.

As a fully functional end product, the Matlab routines used here could be expanded and made comprehensive for all the systems and their components managed by the PPS team. So long as the data were regularly updated, these routines provide fast and accurate information relating to reliability of these systems.

We estimated probabilities via kernel density and found intervals where the ODH and PLC systems are NBU. This study made reasonable assumptions about the distributions of the systems and found the assumptions to be incorrect. We then proceeded with our analysis by deriving our estimates from the data.

Overall we were able to find good agreement with our estimates to the empirical values derived from the data and were able to find where our systems are New Better than Used.

5. ACKNOWLEDGEMENTS

The author would like to thank Mr. Kelly Mahoney at the SNS for his constructive suggestions and direction in preparation of this paper and presenting the problem to the author, and the reviewers for their constructive comments and suggestions for improvement.

REFERENCES

[1] K. Mahoney Manuscript in Preparation. 2014.

[2] N. Levenson *Engineering a Safer World* MIT Press, 2011.

[3] M Hollander and F Proschan *Testing Whether New is Better than Used*, Annals of Mathematical Statistics 43(3) (1972) 1136-1146.

[4] D. Smith *Reliability Maintainability and Risk: Practical Methods for Engineers* Waltham 2011.

[5] P. O'Connor *Practical Reliability Engineering* Wiley, 2002.

[6] M. P. Wand and M. C. Jones *Kernel Smoothing* Chapman and Hall, 1995.

[7] L. Wasserman *All of Nonparametric Statistics* Springer, 2005.

[8] V. A. Epanechnikov *Non-parametric estimation of a multivariate probability density* Theory of Probability and its Applications, 14(1) (1969) 153-158.

[9] www.mathworks.com/matlabcentral/newsreader/view_thread/135924 accessed: October 19, 2014.

[10] M. L. Gamiz, K. B. Kulasekera, N. Limmios and B. H. Lindqvist *Applied Nonparametric Statistics in Reliability* Springer, 2011.

[11] E.L. Kaplan and P. Meier *Nonparametric Estimation from Incomplete Observations* Journal of the American Statistical Association 53(282) (1958) 457-481.

[12] S. Gharramani *Fundamentals of Probability* Pearson, 2005.

[13] W. Wade *Introduction to Analysis* Pearson, 2009.

[14] S. Jiang and D. Kececioglu *Graphical Representations of Two Mixed-Weibull Distributions* IEEE Transactions on Reliability, 41(2) (1992) 241-247.

[15] V. Raykar and R. Duraiswami *Fast Optimal bandwidth selection for kernel density estimation*, Technical Report CS-TR-4774, University of Maryland, College Park, 2005.

[16] J. D. Esary, A. W. Marshall and F. Proschan, *Some Reliability Applications of the Hazard Transform*, SIAM Journal of Applied Mathematics, 18(4) (1970) 849-860.