Towards Performance Prognostics of a Launch Valve

Glenn Shevach[1], Mark Blair[2], James Hing[3], Larry Venetsky[4], Everard Martin[5], John Wheelock[6]

[1,2,3,4]*Naval Air Warfare Center-Aircraft Division Code 4.8.1.4, Lakehurst, NJ, 08733, USA*
glenn.shevach@navy.mil
mark.r.blair@navy.mil
james.hing@navy.mil
larry.venetsky@navy.mil

[5,6]*Naval Air Warfare Center-Aircraft Division Code 4.8.2.1, Lakehurst, NJ, 08733, USA*
everard.martin@navy.mil
john.wheelock@navy.mil

ABSTRACT

Due to its criticality in aircraft carrier steam catapult operations, the performance of the Launch Valve is monitored using timer components to determine the elapsed time for the valve to achieve a set opening distance. Significant degradation in performance can lead to loss in end speed of the catapult and result in loss of aircraft / lives. This paper presents a method of using existing timing data for anomaly detection and predicting when maintenance is required (MIR) for a Launch Valve. Features such as mean and standard deviation of timing values are extracted from clock time data to detect anomalies. Neyman-Pearson Criterion and Sequential Probability Ratio Testing are used to formulate a decision on the degraded state. Once an anomaly is detected, an observation window of the previous *N* filtered samples are used in a risk sensitive particle filter framework. The resulting distribution is used in the prediction of shots until MIR. Performance degradation is extracted from training data and modeled as a third order polynomial. The algorithm was tested on two test sets and validated by Subject Matter Experts (SMEs) supplying the data. An Alpha-Lambda performance metric shows the time predictions until MIR fall inside an acceptable performance cone of 20% error.

1. INTRODUCTION

Steam catapults are among the oldest and most maintenance-intensive systems in the Navy. The steam catapult is a system that launches aircraft from an aircraft carrier by releasing built up steam pressure behind a shuttle

Figure 1: View of Launch Valve in closed and open positions

that pulls the aircraft along the deck. This critical system is largely unchanged from the 1940's – steel, steam and hydraulics that will be with us for the next 40 years. Yet catapults need to perform flawlessly and maintain a system reliability of 99.9999 or the result is loss of aircraft and lives. (Reliability of 99.9 = 140 lost aircraft per year; 99.99 = 14 lost aircraft per year) The Fleet ensures these systems are reliable, but at a very high cost in terms of spares, overhauls and manpower. A reduction in costs could be achieved through prognostic and health management (PHM) methods. The ability to predict impending failures or needed maintenance of these systems in real time, could reduce total ownership costs by decreasing maintenance, inventory, and down time.

The Low Loss Launch Valve (LLLV), hereby known as the Launch Valve, is a hydraulically controlled valve and provides a means for controlling the steam pressure in the catapult power cylinders for launching aircraft (shown in Figure 1). In order to launch the full range of fleet aircraft, the energy of each launch must be tailored for the specific aircraft type and weight, as well as the current wind over deck (WOD) conditions. This is accomplished by adjusting

the opening rate of the Launch Valve to introduce the proper amount of steam. Because of its high reliability requirement, the Launch Valve is designed to have one of the highest operational availabilities compared to all other components within the catapult sub-system. Degradation not being identified quickly can result in additional degradation which could cause a significant loss in end speed and an urgent halt to operations until the degradation was corrected. Insufficient catapult end speed can result in loss of aircraft / lives.

The fleet checks the Launch Valve performance during launching operations with pre-op Blow-Through-No-Loads (BTNL) (no aircraft connected to the catapult shuttle). These times are manually read by an operator, transcribed in a paper log, and typed into electronic spreadsheets hours later. The process is prone to inscription errors. A detailed analysis of Launch Valve performance is manually reviewed upon submission at the conclusion of each month. Subject Matter Experts (SMEs) review clock times to sift out inscription errors and advise for further maintenance actions. This time consuming process relies heavily on the historical knowledge and judgment to decide when a Launch Valve is starting to show signs of degradation. The delay in detailed analysis leaves the potential for degradation to go unnoticed and uncorrected. Continuous real time monitoring of the Launch Valve performance could detect trends in degradation before they reach a critical point.

This paper presents efforts towards the ultimate goal of giving the fleet real time prognostics and health monitoring of the Launch Valve performance during aircraft operations. The algorithm utilizes available Launch Valve clock timing data to detect anomalies and predict when maintenance is required (MIR). Probabilistic techniques are used to detect, with minimum false alarms, the degradation in performance of a Launch Valve and prognostic techniques are used to predict when the degradation will cross a "maintenance needed" threshold. A unique quality to this data is that it is comprised of manually entered time. An operator reads the output of the timers and manually inputs it into a spreadsheet. The algorithm presented takes in timing data over a series of Launch Valve openings that are susceptible to user inscription error.

The paper is structured as follows: Section 2 discusses related works on prognostics and health monitoring of valves. Section 3 provides background information of Launch Valve operation. Section 4 provides the theoretical background for feature extraction, anomaly detection, degradation modeling, and forecasting techniques. Section 5 presents results and discussion using
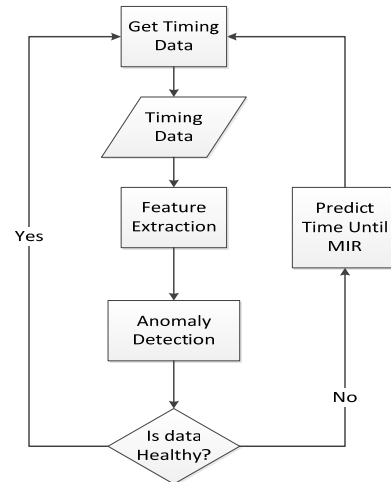

Figure 2: Flow chart for Launch Valve Prognostics

real world Launch Valve timing data and Section 6 concludes the paper with a summary of the findings and future work.

2. RELATED WORKS

Two notable works are related to this paper's efforts. Gomes et. al. developed a health monitoring system for a pneumatic valve using a Probability Integral Transform based technique (Gomes 2010) and Daigle et. al. developed a model-based prognostics approach for pneumatic valves (Daigle 2011). While the Launch Valve in this work is hydraulically controlled, the methods used for pneumatic valve PHM are quite relevant. Diagle et. al. used a Probability Integral Transform to calculate an index of dissimilarity between pressure distributions of monitored and baseline (healthy) valve performance. They were able to use this index of dissimilarity feature to detect increasing degradation and failure of a valve. There was no prediction to failure presented. Timing data of the valve was not utilized. Daigle et. al. constructed a detailed physics-based model of a pneumatic valve that includes models of different damage mechanisms. They use time for the valve to open and close to perform the prognostics. In their work, they focused on the prediction portion of the work and started predictions at pre-defined known points in the historical data where degradation was observed.

3. LAUNCH VALVE OPERATION

The Launch Valve has two (2) clock switches, Clock No.1 and No.2 that are used to measure the time it takes the valve to open 23% and 60% of full open respectively. The beginning portion of launch valve stroke is very dynamic which leads to too much clock time variation in Clock No. 1 to be used as a performance indicator. Clock No. 2 provides less variation in clock times since it measures later in the valve stroke and is therefore used as a performance

indicator. Currently, the Launch Valve performance is monitored by the fleet using Launch Valve Clock No. 2 times from the two daily pre-operational Blow Through No Load (BTNL) launches. The times are compared to limits established in the applicable Maintenance Requirement Card. The fleet conducts both a shot by shot (real time) and long term trend evaluation of the BTNL clock times. NAVAIR Lakehurst also conducts a more detailed analysis of the Launch Valve performance using data (BTNL and aircraft) via the Automated Shot and Recovery Log (ASRL) provided by the fleet.

Degradation in performance of the Launch Valve can be assessed through analysis of this timing data. Performance degradation of the Launch Valve can be caused by increased friction due to loss of lubrication, other internal components providing high friction loads, or parameters outside the normal operating range. Slower clock times are representative of a valve experiencing high internal friction. Faster clock times are representative of a valve leakage in hydraulic fluid downstream. Other factors unrelated to performance are misalignment of the valve and body seat due to surface wear and degraded gasket condition. It can be difficult and costly to install sensors to monitor conditions such as lubrication, wear, gasket condition, etc. This is especially true in these cases where the Launch Valve already exists in a catapult system and cannot be modified. Therefore, a health management solution must be implemented using limited data and feature sets.

4. APPROACH

Figure 2 shows a flow chart for the process that the proposed prognostics algorithm follows.

4.1. Data Preparation

In its current state, the Launch Valve timing data requires some pre-processing by SMEs prior to being fed into the prognostics algorithm. Future work will look to automate the pre-scrubbing process. Raw Clock 2 data contains timing of all launches and blow through no loads. Launches with a low capacity selector valve (CSV) setting have to be identified and removed from the data because CSVs below a specific value do not tend to achieve the Clock 2 switch prior to the "launch complete" signal closing the Launch Valve. This results in inaccurate timing. After this scrub, clock times are compared to existing Clock 2 vs CSV curve baseline (4th order poly fit line) to determine "variation". A 4th order polynomial was found to provide the best fit of the clock times for the range of CSV settings from aircraft operations based on historical data. The next step is the manual review of the data to identify if any shifts in the data occurred signifying a potential shift in the baseline is necessary. Over the life of the catapult the limit switches timing the opening of the Launch Valve will be replaced several times which can cause a shift in the data. If a shift


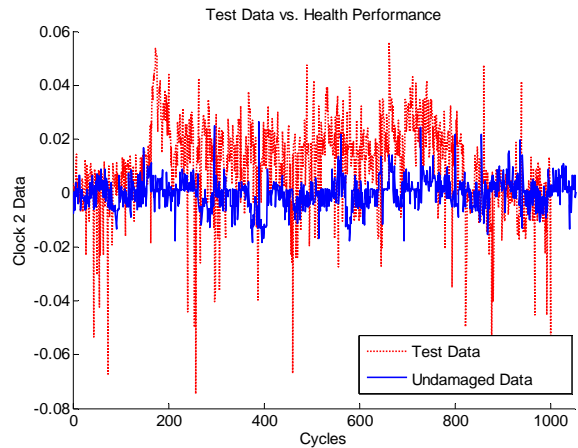Figure 3: Good performance data of opening times of a Launch Valve Over a One Year Period.

Figure 4: Top) Healthy Data (blue solid) vs. Degraded Data (red dashed), Bottom) Gaussian distributions of good performance data (blue solid) and degraded performance high / low. Degraded High means longer clock times than normal, Degraded Low means shorter clock times than normal.

did occur, a new baseline is identified based on identified "good" data. After the baseline is identified, outliers (assumed to be related to inscription errors) are removed based on a +/- 8% variation threshold from the baseline. This helps to eliminate a good portion of transcription errors but not all.

The data used in this study was broken into training sets of known Launch Valve performance data and two test sets of unknown performance data (known by SME supplying the data). Specifically, the training sets contained 27,622 sequential shots of healthy performance data and 11,882 sequential shots that contained degraded performance within the data set. Test Set 1 contained 19,355 sequential shots and Test Set 2 contained 10,648 sequential shots.

4.2. Feature Extraction

The prognostics algorithm presented in this work, starts with the assumption that the following data has been

received: shot number, Clock 2 times, and base line times for all catapult shots. The extraction of these times was described in the previous section. To account for any shift in the Clock 2 timing of the valve, the clock times ($time_{obs}$), are normalized using the baseline time ($time_{exp}$), resulting in **Clock 2 Data** as illustrated in Eq. (1).

$$\textbf{Clock 2 Data} = \frac{time_{obs} - time_{exp}}{time_{exp}} \qquad (1)$$

The algorithm tracks all aircraft shots. Both BTNLs and aircraft shots are used to track performance. Figure 3 shows an example set of Clock 2 data of a healthy Launch Valve over a one year period.

The distribution of the Clock 2 data, C, over N launch cycles, $p_N(C)$, data tends to fit a Gaussian distribution of the following form:

$$p_N(C) \rightarrow \frac{1}{\sigma\sqrt{2N\pi}} e^{-\frac{(C-N\mu)^2}{2N\sigma^2}} \qquad (2)$$

which is the formula for a Gaussian distribution with mean $N\mu$ and variance $N\sigma^2$.

Based on consultations with SMEs, it was determined that degraded operation resulted in a shift of the mean and a change in the standard deviation of the clock times. There are two different degraded modes. Data that has an increasing mean (slower clock times, Degraded High) can be representative of a valve experiencing high internal friction; while data that has a decreasing mean (faster clock times, Degraded Low) can be representative of a valve leakage in hydraulic fluid downstream. An example of this is demonstrated in Figure 4 where the blue data (solid line) represents a healthy Launch Valve and the red data (dots/dashes) represents a valve operating in a degraded condition (low – dashed line, high – dotted line). These distribution functions were extracted by analyzing the training set of known healthy, degraded low, and degraded high valve performance data. The mean and standard deviation are used as features to detect anomalies in the clock data.
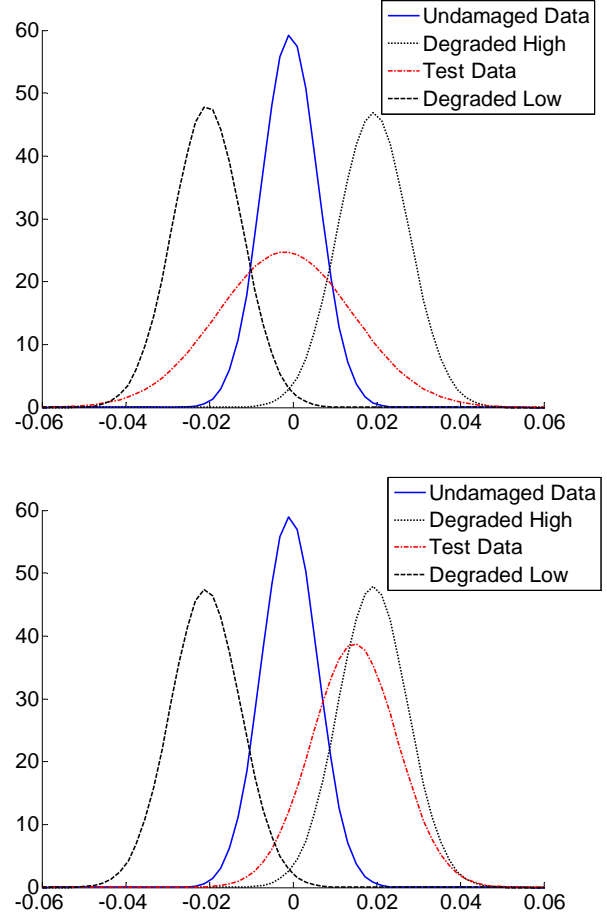

Figure 5: PDFs of performance data. Top) Test data is still in the good performance range. Bottom) Test data has shifted into the degraded high range.

4.3. Anomaly Detection

This work implements a data driven approach for detection of degradation in Launch Valve performance. The problem simplifies to an anomaly detection problem, i.e. detecting when the incoming signal (features) are diverging from a historically estimated healthy state. Parameters for the healthy state are extracted from a known healthy training set of data and used in the comparison against incoming data. A hypothesis test is conducted using the Neyman-Pearson Criterion (Lehmann 1986). Neyman-Pearson is a probabilistic method used to classify data points in a null or alternative hypothesis by calculating a likelihood ratio and comparing it to a threshold.

In the case of the Launch Valve, the two different degraded modes lead to two alternative hypotheses, Degraded High or Degraded Low. Table 1 shows the designation of these states.

Table 1: Neyman-Pearson Hypotheses

H_0	Null Hypothesis that the Launch Valve is healthy
$H_{1_{High}}$	Hypothesis that the Launch Valve is degraded indicated by slower clock times
$H_{1_{Low}}$	Hypothesis that the Launch Valve is degraded indicated by faster clock times

Figure 5 Top provides a visual representation of the various performance distributions. The black probability distribution functions (PDFs) represent degraded low and high data, the blue PDF is an undamaged set of data, and the red PDF is an example set of test data. The increased standard deviation in the test data may be due to intermittent inconsistencies in lubrication during operation. The Neyman-Pearson Criterion calculates the likelihood ratio, $L(x)$ (shown in Eq. 3), which is the ratio of the probability of a data set belonging to the alternative hypothesis versus the null hypothesis. The probability of accepting $H_{1_{High}}$ increases when the test dataset starts to shift, as seen in Figure 5 Bottom.

NOTE: For future reference, any degraded state will be represented by H_1 unless a low/high degraded state is specifically stated.

Two false alarm rates, Type I Error and Type II Error, must be specified to correctly classify an anomaly. Table 2 below shows the designations of both of these errors.

Table 2: False Alarm Rate Designation

P_{FA_I}	Probability of Type I Error (False Positive: Conclude damage is present falsely)
$P_{FA_{II}}$	Probability of Type II Error (False Negative: Conclude damage is not present falsely)

The probability of a Type I Error was set to 0.01 yielding a probability of detection of 99%. The likelihood ratio is then calculated to help classify when the measured data set **x** signifies degraded operation. If this ratio is greater than one, there is a higher probability of accepting the alternative hypothesis.

$$L(\boldsymbol{x}) = \frac{p(\boldsymbol{x}|H_1)}{p(\boldsymbol{x}|H_0)} \quad (3)$$

To better utilize the measurement distribution, a window (size W=100 launch cycles) of timing data $\boldsymbol{x}$ is used in the likelihood ratio as follows:

$$L(\boldsymbol{x}) = \frac{\prod_{i=1}^{W} p(x_i|H_1)}{\prod_{i=1}^{W} p(x_i|H_0)} \quad (4)$$

The next phase of anomaly detection implements a Sequential Probability Ratio Test (SPRT). The SPRT evaluates deviations of the actual signal from the expected signal (healthy data) based on distributions instead of a single threshold value to determine if data belongs to a degraded state. SPRT uses the log of the likelihood, L(x), in a sequential analysis. (Wald, 1947). The cumulative log-likelihood is calculated, as seen in Eq. (5), and compared against lower and upper thresholds a and b to determine the next course of action (Table 3). As a new sample becomes available, the observation window shifts, calculating a new likelihood ratio and SPRT value.

$$SPRT_i = SPRT_{i-1} + \log\big(L(\boldsymbol{x_i})\big) \quad (5)$$

Table 3: SPRT Comparison Statements

$a < SPRT_i < b$	Continue monitoring
$SPRT_i \geq b$	Accept H_1
$SPRT_i \leq a$	Accept H_0

With a set probability of 1% for a Type I Error and a set probability of 5% for a Type II Error, thresholds a and b are calculated using Eq. (6) and Eq. (7) respectively.

$$a = \ln\left(\frac{P_{FA_{II}}}{1 - P_{FA_I}}\right) = -2.99 \quad (6)$$

$$b = \ln\left(\frac{1 - P_{FA_{II}}}{P_{FA_I}}\right) = 4.55 \quad (7)$$

H$_1$ is accepted when the SPRT calculation exceeds the b threshold. This concludes there is enough data to support the decision to determine an anomaly has been detected. The SPRT is then reset if the value has declined consecutively for 20 iterations. If H_0 is accepted, the cumulative log-likelihood ($SPRT_i$) is reset to zero to restore sensitivity to small changes in degradation. A similar approach to anomaly detection was implemented by Cheng et. al. for monitoring environmental and operational stress profiles of robotic vehicles (Cheng, 2008).

4.4. Degradation Model

A third order polynomial was chosen as a data-driven damage progression model based on a best fit of multiple degradation sections from the training sets. SMEs also helped to define the ranges for initial parameter distributions for the model parameters based on their experiences with historical performance degradation trends. The performance degradation model follows Eq. (8) where $\boldsymbol{a}$, $\boldsymbol{b}$, and $\boldsymbol{c}$ are model coefficients, $\boldsymbol{T}$ is the translation parameter allowing the model to adapt to shifting states of degradation, $\boldsymbol{y}$ is the degraded state prediction of the next shot, i is the sample index (with index 1 being the detected start of degradation), and dt is the cycle increment which was set to 1 (each shot increments by 1).

$$\boldsymbol{y}_i = \boldsymbol{a}(i + dt + \boldsymbol{T})^3 + \boldsymbol{b}(i + dt + \boldsymbol{T})^2 + \boldsymbol{c}(i + dt + \boldsymbol{T}) \quad (8)$$

Parameters **a**, **b**, **c**, and **T** are initialized after an anomaly is detected and are updated via the particle filter (described in the next section) for as long as the data classifies the Launch Valve operation as degraded.

The effect of loading conditions (varying aircraft weights) on the degradation of the launch valve performance is negligible. The CSV controls the launch valve rate of opening regardless of what aircraft is on the catapult. In other words, regardless of the aircraft type, if a value of CSV 200 is used to launch a F/A-18 or an EA-6B aircraft (two different weight aircraft), the launch valve clock time should be the same.

4.5. Prediction

Once an anomaly is detected, a particle filtering (PF) based prognostic algorithm takes over. PF prognostic algorithms have become a common method in the state of the art prognostics. A PF is used to provide estimations of distributions of model parameters using a window of observations. This is accomplished using Bayesian inference, based on Bayes' Theorem as seen in Eq. (9), where Θ is a vector of unknown parameters (a,b,c,T), $p(\Theta)$ is the prior PDF of these parameters, z is the vector of observed data (clock 2 time), $p(\Theta|z)$ is the posterior PDF of Θ conditional on z and $L(z|\Theta)$ is the likelihood of the observed data given the parameters (An, 2012).

$$p(\Theta|z) \propto L(z|\Theta)p(\Theta) \qquad (9)$$

The particle filter utilizes a sequential method of passing prior estimations into the current step to produce the estimations for the next step. In particular, this work implements a simplified version of the Risk-Sensitive

Particle Filter (RSPF) presented by Orchard et. al. (Orchard, 2010). The RSPF maintains a subset of particles in the high-risk, low-likelihood realm to maintain coverage in these areas when incoming data causes convergence of particles to a single particle or narrow distribution. In this work, twenty percent of the particles are allocated to maintain distribution within the risk sensitive areas.

Input into the PF is timing data that has been filtered with two passes of an exponential moving average filter (EMAF) as shown in Eq. (10). Development with training data supported using parameters $\alpha = 0.003$ on the first pass and 0.03 on the second pass. The EMAF is an infinite impulse response discrete filter that provides low latency.

$$EMAF_i = \alpha f_i + (1 - \alpha)EMAF_{i-1} \qquad (10)$$

The degradation model parameters are estimated using a 10 sample window of EMAF data. Using a sample from the EMAF data, a likelihood calculation is performed and 1000 particle weights are updated. Each particle represents a particular parameter configuration with a particle weight
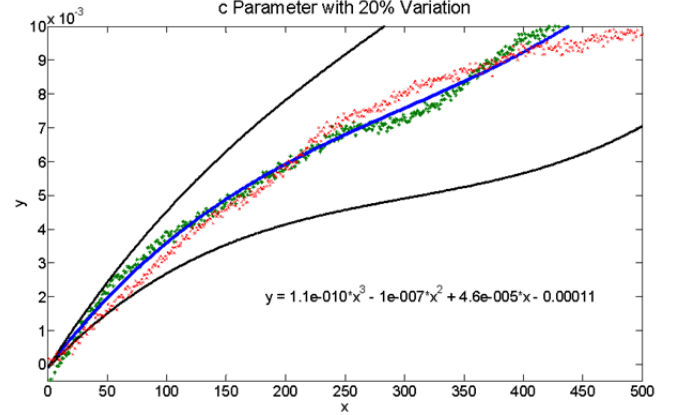

Figure 6: Performance degradation plots. Two examples showed (darker dots, lighter dots). Third order model fit to data. 20% Bounds on c parameter shown by black lines.


Figure 7: Top) Test set 1. The algorithm classified this test set as containing all healthy data, Bottom) Test set 2. The algorithm classified this test set as containing degraded performance data.

based on its likelihood. These weights are then used in the likelihood calculation for the next measurement sample of the current EMAF window. Parameters are updated for each sample of the window and the resulting particle weights are used in a third order model to generate each particle prediction.

Once predictions have exceeded the failure threshold (defined by the SME), each particle contributes to the time until MIR PDF. When a new measurement data point is acquired, the EMAF output is updated and the particle filter

window is shifted. The shifted window is then passed through the process to update the parameter weights and provide a prognosis, utilizing a portion of the weights from the previous measurement. The prognosis process repeats, resulting in updated MIR predictions as the degradation progresses.

5. RESULTS

The algorithm was tested against two sets of data, shown in Figure 7, of unknown classification to the program (but known by the SME who supplied the test sets). For each classification test, the algorithm was fed the test data cycle by cycle, as if it was being deployed in real time. Once the observation window is filled, each data point was classified as belonging to a degraded state or a healthy state. Overall the test sets were classified as "healthy" if they had no anomaly detections and "degraded" if anomalies were detected. The algorithm classified Test Set 1 as containing only healthy data and Test Set 2 as containing degraded performance data. The SME validated that this was the correct classification for the data that he supplied. Furthermore, for Test Set 2, the algorithm identified locations in time for which degraded performance was identified (shown in Figure 8).

At the start of identified degradation (rising edge on plot in Figure 8), the prediction algorithm took over and predicted out when the performance data would cross a pre-defined "maintenance needed" threshold. An example is shown in Figure 9 where an anomaly was detected around cycle shot 4290 and predictions were made for the remaining cycles until maintenance would be required. The figure shows an example of predictions to MIR at about 50% remaining time until MIR.

To assess the quality of the prediction for Test Set 2 (shown in Figure 9), the Alpha-Lambda performance metric is used (Saxena 2009). The Alpha-Lambda performance metric is an off-line metric that determines whether the prediction falls within the specified levels of a performance measure at particular times. The time instances are specified as a percentage of total remaining life (cycles until MIR in this case) from the point the first prediction is made. Accuracy, defined as the prediction accuracy of cycles until MIR, is set to be alpha*100% of the actual cycles until MIR. In this case, an alpha of 0.2 was used. Results from Test Set 2 consistently showed the prediction of remaining cycles until MIR fell within the 20% accuracy (alpha = 0.2) with approximately 70% (lambda = 0.7) of the remaining cycles until MIR remaining. This can be seen in Figure 10. Early predictions in the normalized prognostic window tend to fluctuate outside the Alpha-Lambda cone due to wide spread in the distribution of particles used in the particle filter. As more degraded data is acquired, the particle distribution tightens as the particle filter begins to converge on a particular degradation model.
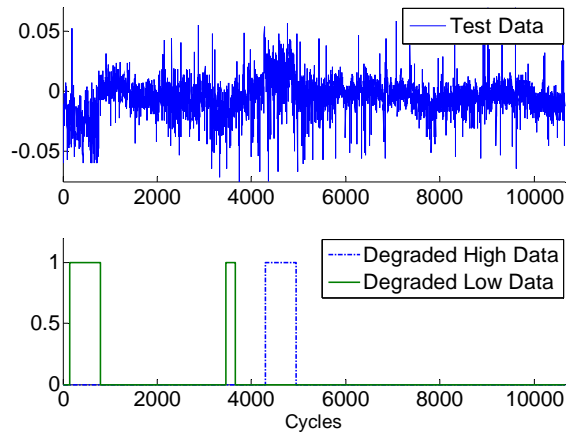

Figure 8: Test set 2 with algorithm identified locations with degraded performance in both low (green) and high (blue) levels. "Low" means timing is shorter than normal, "High" means timing is longer than normal.
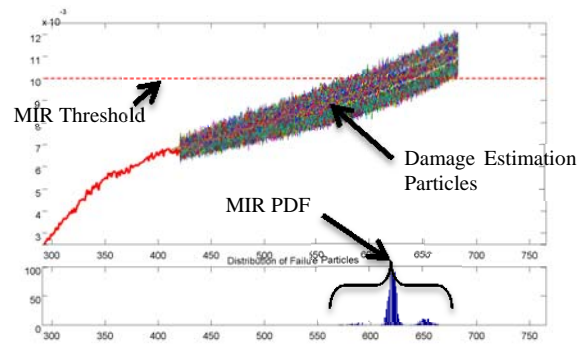

Figure 9: Particle Filter Estimation of degradation and MIR PDFs.

Figure 10: Alpha-Lambda Performance with 20% error bound. Prediction until MIR showing Median, 5%, and 95% confidence levels (Cl).

6. CONCLUSIONS AND FUTURE WORK

For the Low Loss Launch Valve, the method of extracting and using features from timing data such, as mean and standard deviation, to detect anomalies using Neyman-Pearson Theorem and SPRT has been shown in the previous section to produce promising results. The prediction of the remaining time until MIR with a risk sensitive particle filter using a third order model has also been shown to produce results within an acceptable accuracy window. This is a step towards allowing the Launch Valve performance analysis to be handled automatically in real-time onboard ship and provide timely status information to the fleet.

The next step toward achieving an automated PHM solution for the Low Loss Launch Valve is to automate the process of pre-scrubbing the data which is currently handled by the SME. The automated pre-scrub would need to receive raw

clock timer information (CSV setting and Clock 2 time), screen out low CSV launches not useable for review, and properly identify baseline shifts without input from users. The algorithm needs to handle varying levels of noise / error in the data, much due to transcription errors. It is possible that future upgrades to the launch system could incorporate added sensors and electronic logging to automatically record the timing data, thereby eliminating transcription error issues.

Acquiring more test data sets would further verify / validate the PHM methodology presented in this work. With more data, it is possible that supervised learning algorithms such as neural networks could be used to improve upon classification methods and anomaly detection. Future work will also include methods of identifying healthy data in real-time data sets (deployed system) and use that to set anomaly detection and prognostics parameters. This would reduce reliance on fleet historical data and would tailor PHM methods to each specific Launch Valve system through its life span.

REFERENCES

An, D., Choi, J., Kim, N.H., (2012), "A Tutorial for Model-based Prognostics Algorithms based on Matlab Code," *Annual Conference of Prognostics and Health Management Society*, vol. 3, no. 122, September 2012.

Cheng, S. (2008), Autonomous Prognostic Monitoring Device. *Proceeding for 62nd Meeting of the Society for Machinery Failure Prevention Technology (MFPT)*, Virginia Beach, VA, May 2008

Daigle, M., Goebel, K. (2011), "A model-based prognostics approach applied to pneumatic valves," *International Journal of the PHM Society*, vol. 2, no. 2, pp. 1-16.

Gomes, J., Ferreira, B., Cabral, D., Glavao, R., Yoneyama, T. (2010), "Health Monitoring of a Pneumatic Valve Using a PIT Based Technique," *Annual Conference of Prognostics and Health Management Society*

E. Lehmann (1986), *Testing Statistical Hypotheses*, New York: Wiley, 1986.

Orchard, M., Tang, L., Saha, B., Goebel, K., Vachtsevanos, G. (2010), "Risk-Sensitive Particle-Filtering-based Prognosis Framework for Estimation of Remaining Useful Life in Energy Storage Devices," *Studies in Informatics and Control*, vol. 19, no. 3, September 2010.

Saxena, A; Celaya, J.; Saha, B.; Saha, S.; Goebel, K. (2009), "Evaluating Algorithm Performance Metrics Tailored for Prognostics," *IEEE Aerospace Conference*, vol., no., pp.1,13, 7-14 March 2009.

Wald, A. (1947), "Sequential Analysis", *John Wiley & Sons*, New York, NY,.

BIOGRAPHIES

Glenn Shevach is an Electrical Engineer in the Advanced Technology Projects Branch Code 4.8.1.4. at NAWCAD in Lakehurst, NJ. He holds a Masters Degree in Electrical Engineering from Stevens Institute of Technology (2010). He is the Advanced Arresting Gear S&T / Prognostics Lead.

Mark Blair is a Computer Engineer in the Advanced Technology Projects Branch Code 4.8.1.4. at NAWCAD in Lakehurst, NJ. He is currently working towards a Masters Degree in Systems Engineering from Stevens Institute. His current work focuses on S&T related topics for Aircraft Launch Recovery and Support Equipment (ALRE & SE).

James Hing is a Mechanical Engineer in the Advanced Technology Projects Branch Code 4.8.1.4 at NAWCAD in Lakehurst, NJ. He holds a PhD in Mechanical Engineering from Drexel University (2010). His current work focuses on robotics and prognostics solutions for ALRE & SE applications.

Larry Venetsky is an Electrical Engineer in the Advanced Technologies Project Branch Code 4.8.1.4 at NAWCAD in Lakehurst, NJ. He holds MS in Electrical Engineering from Drexel University (1993). His current work focuses on robotics and machine learning for ALRE & SE applications.

Everard Martin is a Mechanical Engineer in the Steam Catapult Launcher In-Service Branch, Code 4.8.2.1 at NAWCAD Lakehurst, NJ. He holds a BSE in Mechanical Engineering from Drexel University (1981). He has been working Steam Catapult Performance and Certification since June 1989.

John Wheelock is a Mechanical Engineer in the Steam Catapult Launcher In-Service Branch, Code 4.8.2.1 at NAWCAD Lakehurst, NJ. He holds a BSE in Mechanical Engineering from Arizona State University (2002). He has been working Steam Catapult Performance and Certification with a primary focus on catapult performance monitoring since July, 2002.