Estimation of Life Consumption for Advanced Drilling Tools

Dustin Garvey[1], Martin John[1], and Jörg Baumann[1]

[1] *Baker Hughes, Celle, Germany*

dustin.garvey@bakerhughes.com
martin.john@bakerhughes.com
joerg.baumann@bakerhughes.com

ABSTRACT

Prognostics has the potential to be very valuable in many industries. This is especially the case in the petroleum industry where the costs of tool failure are extremely high and continue to increase. Previous attempts have been made to predict the remaining useful life of drilling tools. While the developed methods were shown to be able to accurately predict the remaining useful life, the data requirement was such that they had limited or no viability in "real world" operations. This paper builds on previous work in this area by developing a new life consumption estimation model that has been specifically designed to ensure that it can be viable in the "real world". The developed model was shown to be able to estimate the life consumed of an individual drilling tool to within 4-12% with uncertainties of ±15-35%.*

1. INTRODUCTION

In the petroleum industry, drilling tool reliability is critical for performance. With rig costs reaching $350,000 a day (Wand et al., 2006) and increasing drilling profile difficulty, there is an ever-increasing need for more reliable tools. Furthermore, recent research by Brehme and Travis (2008) has found that tool unreliability has produced tool failure rates on the order of 33%, which increases the urgency for the development of new processes and methods to support reliability improvement.

One way to improve reliability is to develop a prognostic model that is able to accurately estimate tool weariness such that it can be removed from service prior to failure. Furthermore, to be useful in the drilling industry, a method needs to be developed that readily fits into established operational practices such that it has a higher chance of being realistically used in day-to-day operations. This paper describes and demonstrates a new life consumption estimation algorithm that meets these requirements. This paper will:

- introduce advanced drilling systems,
- describe the life consumption model, and
- present results of a feasibility study.

2. ADVANCED DRILLING SYSTEMS

Modern drilling systems can be thought of as being composed of two major components:

1. drill pipe and the
2. bottom hole assembly (BHA).

The function of the drill pipe is to maintain hole integrity and connect the surface to the BHA. The BHA is the “heart” of the drilling system and is a collection of specialized tools that provide different sets of functionality to enhance the performance and efficiency of the drilling process. In a typical directional drilling application, the BHA is composed of:

- drill bit,
- steering tool,
- power generation tool,
- communication tool, and
- formation evaluation tool(s).

At this point, it is important to note that modern BHAs are massive systems and are exposed to extremely stressful environments. As depicted in Figure 1, modern BHAs:

- have lengths on the order of 50 ft,

- can drill to depths more than seven miles below the earth's surface (ExxonMobil 2007), and
- are capable of producing highly complex well profiles.

One of the most critical tools is the steering system (Wand et al., 2006; Brehme and Travis, 2008). For this reason, it was selected as the first target for the development of a prognostic system. In the present application, the steering system is composed of three hydraulic ribs that extend to "push" the BHA in the desired direction.

Current practices have tool components that are deployed in different wells across its life between maintenance. Keeping track of data for an individual tool and/or its components over its life is very difficult. A method is needed that allows us to make an assessment of consumed life for an arbitrary run in that device's history. This way, we only need to carry a single number along with the device and thereby make tool weariness assessment practical.

3. LIFE CONSUMPTION ESTIMATION

This work is not the first step to try to develop a prognostic model for drilling tools (Hines and Garvey, 2008; Garvey et al. 2009). The model that was developed was referred to as the path classification and estimation (PACE) model. While the PACE was demonstrated to be able to produce accurate remaining useful life (RUL) estimates, several factors limited its applicability to deployed tools. This section will describe the initial PACE model, point out its short comings, and then discuss how it was modified to address its problems.

3.1 Initial PACE Model

The initial PACE model can be most easily understood by considering the problem outlined in Figure 2. Here, we have functional approximations of example cumulative stresses, $U_i(t)$, for devices until failure at time T_i. We also have the observed stress of another device at time t^*, denoted by $u(t^*)$.

In the original formulation of the PACE, the remaining useful life of the new device is estimated according to the following three step process:

1. calculate expected stresses and RULs,
2. classify the stress according to expected values, and
3. estimate the RUL.

Let's now step through this process using the example presented in Figure 2.

The first step in the process is to calculate expected cumulative stresses and RULs had the device been progressing along each of the four example paths. To obtain the expected stresses, we evaluate the functional approximations at time t^*. To obtain the expected RULs for each of the example paths, we subtract the time t^* from each of the failure times. The results of these operations are the following vectors.

$$\mathbf{f}(t^*,\mathbf{\Theta}) = \begin{bmatrix} f_1(t^*,\mathbf{\theta}_1) \\ f_2(t^*,\mathbf{\theta}_2) \\ f_3(t^*,\mathbf{\theta}_3) \\ f_4(t^*,\mathbf{\theta}_4) \end{bmatrix} \quad \mathbf{L}(t^*) = \begin{bmatrix} T_1 - t^* \\ T_2 - t^* \\ T_3 - t^* \\ T_4 - t^* \end{bmatrix} \tag{1}$$

Here, $f_i(t^*, \mathbf{\theta}_i)$ is the functional approximation, $\mathbf{\theta}_i$ are the regressed parameters, and T_i is the observed failure time of the i^{th} path.

What these vectors tell us is that had we been progressing along each of the example paths, then at time t^*, we would expect to have the calculated stresses and remaining lifetimes.

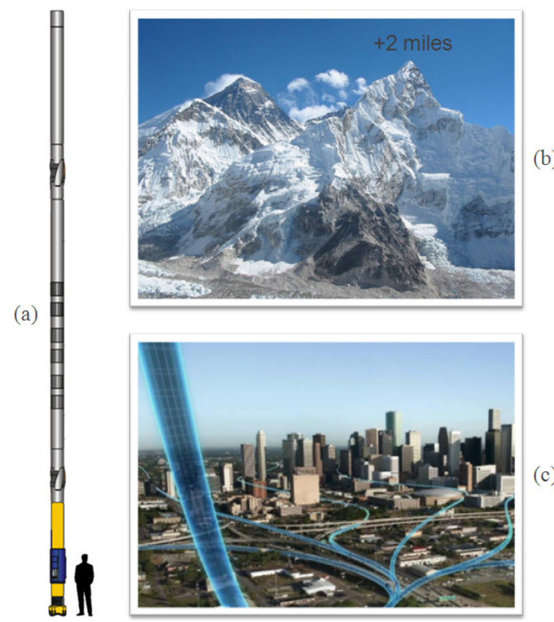

Figure 1. Modern BHAs (a) have typical lengths of 50 ft, (b) can drill to depths greater than the height of Mt. Everest, and (c) are able to produce complex well geometries comparable to modern cities such as Houston, TX.

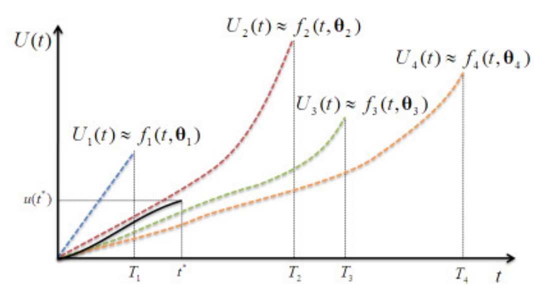

Figure 2. Functional approximations and failure times of historical and example cumulative device stresses.

Now that we have exemplar stresses and lifetimes, we are ready to classify the observed stress according to the examples. To accomplish this, the Gaussian kernel function is used (Fan and Gijbels, 1996):

$$w_i = \frac{1}{\sqrt{2\pi h^2}} e^{-d_i^2 / 2h^2} \quad (2)$$

Here, d_i is the Euclidean distance of the observed stress from the i^{th} expected stress:

$$d_i = \sqrt{\left[f_i(t^*, \boldsymbol{\theta}_i) - u(t^*)\right]^2} \quad (3)$$

Additionally, h is the kernel bandwidth, which controls how close the stress needs to be to the exemplar stress to be deemed similar and therefore receive large weights, w_i. At this point, we have a vector of weights, $\mathbf{w}$, whose elements indicate how similar the observed stress is to the expected stresses.

The final step in the initial PACE model is to use the results of the classification to estimate the RUL. This is accomplished by calculating a weighted average of the expected RULs where the weights are the similarities of the observed stress, $u(t^*)$, to the expected stresses. For our example, the RUL can be estimated according to the following equation.

$$\hat{l} = \frac{w_1(T_1 - t^*) + w_2(T_2 - t^*) + w_3(T_3 - t^*) + w_4(T_4 - t^*)}{w_1 + w_2 + w_3 + w_4} \quad (4)$$

3.2 Shortcomings of the Initial PACE Model

While the initial PACE was demonstrated to produce accurate RUL estimates, its incorporation into daily operations is limited. The reason for this is that in order to calculate the RUL, we need to know the cumulative stress that a particular device has been exposed to for a given time in its history. In other words, we need all of the stress data for a particular tool beginning after the last significant maintenance.

At first glance, this requirement doesn't seem that demanding, but for current operations the infrastructure to autonomously capture and log this data is not in place. Furthermore, even if the infrastructure was in place, the data requirement significantly increases the complexity of the analysis and thereby limits its use in the "real world".

For a prognostic algorithm to be viable, it needs to be able to analyze segments of data and produce a metric that can be easily communicated and compounded to assess the health of a device. To accomplish this goal, the initial PACE was adapted to analyze a particular segment of data to produce a life consumption estimate. This life consumption estimate can then be easily passed along and compounded to calculate the total life consumption of the tool at any given time. In its adapted form, to estimate the total life consumed of a particular device only the stress data from the latest deployment and the running life consumption total is needed. By simplifying the data requirements, we are able to produce a model that is more viable in "real world" operations.

3.3 Modified PACE Model

To begin, we are going to shift our previously used problem statement. Rather than ask the question, “Can we estimate the remaining useful life?” We’re going to ask, “Can we estimate the consumed life?” The reason that we make this shift is that if we can estimate the consumed life in an individual run, then the process of estimating the total consumed life is simply a matter of accumulating the consumed lives of the previous runs. In other words, to assess the health of an individual tool we only need to carry around a handful of numbers and not the stress data for the individual tool. Let’s now take a look at how we are going to do this.

While we’re asking a fundamentally different question, we can still use the PACE framework. Mainly, we are still going to classify the path and then use the result of this classification to make an estimate of the consumed life. The main difference is that instead of calculating the expected stresses and expected remaining lives, we calculate the expected life consumptions and the observed "shape" of the stress accumulation. Let's step through our hypothetical example once more, using the modified PACE model.

The first step in the process is to create exemplar memory vectors of the "shapes" of the stress accumulation and the corresponding life consumption fractions at time t^*. To create the vector of "shapes", we simply concatenate the regression parameters of the functional approximations. To create the vector of expected life consumptions, we simply divide the current time t^* by each of the failure times.

$$\boldsymbol{\Theta} = \begin{bmatrix} \boldsymbol{\theta}_1 \\ \boldsymbol{\theta}_2 \\ \boldsymbol{\theta}_3 \\ \boldsymbol{\theta}_4 \end{bmatrix} \qquad \mathbf{L}_c = \begin{bmatrix} t^*/T_1 \\ t^*/T_2 \\ t^*/T_3 \\ t^*/T_4 \end{bmatrix} \quad (5)$$

What these vectors tell us, is that had we been absorbing stress at a rate that is characterized by the example history's shape, we would expect to have consumed t^* / T_i of its life at time t^*.

Before we can classify the current stress accumulation according to the expected values, we need to characterize the accumulation given its history for the current use. For the current use, let's suppose that we have the following sequence of observations:

$$\{\ u(t_1), u(t_2), u(t_3), \ldots, u(t^*)\ \} \quad (6)$$

We can then use regression to approximate the above sequence of observations by a function.

$$u(t) \approx f(t, \hat{\boldsymbol{\theta}}) \quad (7)$$

We are now ready to classify the device stress accumulation according to the expected accumulations. To do this we use Eq. 2, with the exception that the distance is now the Euclidean distance of the regressed parameter to the expected parameter had the device been accumulating stress in a similar manner as the example history. What the distances tell us, are how different the "shape" of the stress accumulation is from the examples.

$$d_i = \sqrt{\left(\theta_{i,1} - \hat{\theta}_1\right)^2 + \cdots + \left(\theta_{i,p} - \hat{\theta}_p\right)^2} \qquad (8)$$

Here, p is the number of parameters in the regressed function. As with the initial PACE, the distances are converted to similarities or weights via Eq. 2.

We can finally calculate the life consumption estimate by calculating a weighted average of the expected life consumptions using the similarities of the accumulation to the expected accumulations as weighting parameters.

$$\hat{l}_c = \frac{w_1\left(t^*/T_1\right) + w_2\left(t^*/T_2\right) + w_3\left(t^*/T_3\right) + w_4\left(t^*/T_4\right)}{w_1 + w_2 + w_3 + w_4} \qquad (9)$$

While the structure of the modified PACE is not significantly different than its original incarnation, it does tackle the prognostics problem in such a way that should increase its viability in "real world" operations. The next section will demonstrate the modified PACE using data collected from tools that were operated worldwide.

4. RESULTS

Operational data was collected from a rotary steering system (RSS) for use in this study. The RSS was selected since it has the largest impact on BHA reliability (Wand et al., 2006). Lateral vibration data was collected from the operational data between the last significant maintenance until tool failure. The vibration data was used to create stress paths by calculating a running sum of the vibration.

Before continuing, let's take a look at an example vibration signal and see how it is used to create a cumulative stress path. An example of the lateral vibration over the entire history of a tool is presented in the top chart of Figure 3. The cumulative sum of the lateral vibration is then calculated to generate the accumulated stress path presented in the lower plot of Figure 3. Notice that the lateral vibration hovers consistently at values around 2 Gs (sharp drops are instances where drilling is halted). This translates to a linear stress path over the history of the tool. The slope of the path will be seen to be very important in characterizing stress accumulation.

Before the modified PACE model was trained, vibration data was collected from tools beginning at the time at which they received a specific level of maintenance until they either received additional maintenance or failed. The resulting stress paths collected from these runs are presented in Figure 4. For this work, 262 histories of unfailed (blue) and 21 histories of failed (red) tools were created.

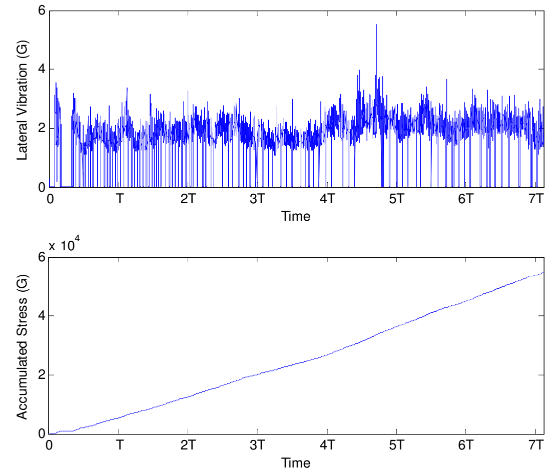

Figure 3. Example of lateral vibration data (top) and calculated accumulated stresses (bottom).

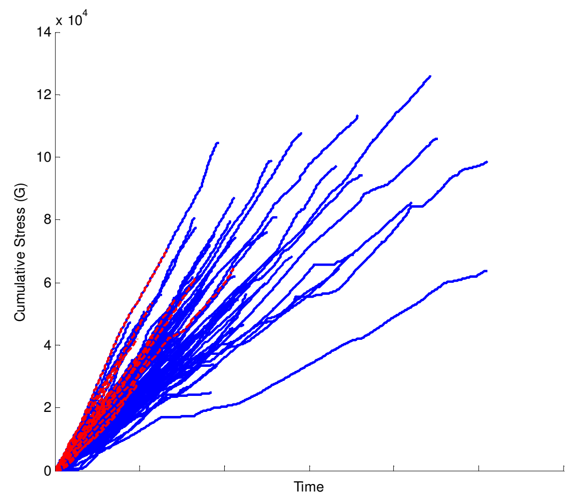

Figure 4. Cumulative stress paths for unfailed (blue) and failed (red) tools.

One feature that is readily apparent from the curve data is that there is a consistent lumping of the failed tools along paths with slopes that high relative to most of the unfailed paths. This is corroborated by the fitted statistical distributions of the slopes presented in Figure 5. While the distributions are not completely separable, we still have a consistent indication that the rate at which the tool absorbs vibration is helpful for inferring its useful life.

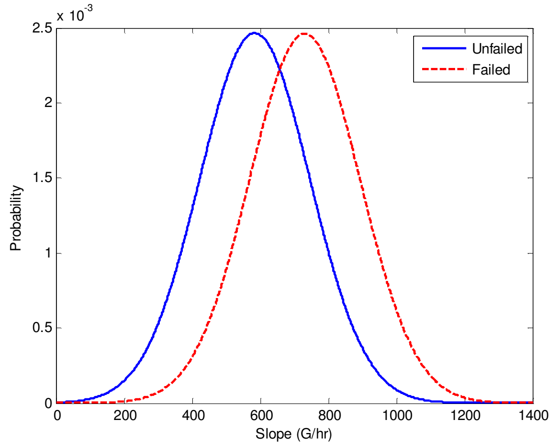

Figure 5. Fitted statistical distributions of unfailed (blue) and failed (red) slopes of the cumulative stress path.

The modified PACE model was trained on a sample of 18 of the failed paths and 3 exemplary unfailed paths. The reason for the use of unfailed paths in the training set can be seen in Figure 4. Here, we can see that the failed paths have a consistently high slope and therefore do not completely bound the overall data space. In other words, the PACE needs to have the capability to estimate the life consumption when the slopes fall outside the range of the failed runs. Furthermore, the unfailed paths were selected on the basis that the tool had a lifetime that exceeded expectations (2 examples) or had been exposed to excessive (i.e. largest total) vibration without failure (1 example).

To test the modified PACE model, it was used to predict the life consumptions for the 3 failed paths not included in the training set. The accuracy was measured in terms of how closely the life consumption estimates met the target. The uncertainty was estimated by analyzing the variance and biases of the life consumption estimates for different random samples of the training paths.

Before discussing the results, it is important to take a step back and briefly discuss the uncertainty analysis technique implemented in this work. The Monte Carlo based uncertainty analysis technique discussed in Hines et al. (2008) is used. For this work, 50 individual PACE models are created by bootstrap sampling (Efron and Tibshirani, 1994) the training paths. Next, these models were used to estimate the life consumptions of the test paths. The 95% confidence interval of the i^{th} life consumption estimate can be approximated by the following equation:

$$\hat{l}_{c,i} \pm 2\sqrt{Var(\hat{l}_{c,i}) + [Bias(\hat{l}_{c,i})]^2} \qquad (10)$$

The variance was calculated by simply taking the variance of the 50 model estimates for the i^{th} observation. The bias was calculated via the bias-variance decomposition of the mean squared error (Tamhane and Dunlop, 2000).

$$Bias(\hat{l}_{c,i}) = \sqrt{MSE(\hat{l}_{c,i}) - Var(\hat{l}_{c,i})} \qquad (11)$$

The results of the modified PACE model are presented in Table 1. In the table, the accuracy of the model is quantified by the mean absolute error (MAE) and the uncertainty is quantified by the mean uncertainty over each of the test estimates. Notice that while the model is able to predict the life consumption with an accuracy on the order of 90%, the uncertainties are quite large. This result is expected since we are training on relatively few examples. As additional histories are collected, the uncertainties should decrease. Furthermore, the accuracies should continue to improve.

Table 1. Modified PACE Model Accuracy and Uncertainty Test Metrics

Test Path	MAE	Mean Uncertainty
1	11%	±27%
2	4%	±15%
3	12%	±35%

To help visualize these results, refer to the charts of the life consumption estimates and their corresponding uncertainties presented in Figure 6 for the three test paths. Notice that for all three test paths the estimates (solid black line) track the actual values (dashed black line) to a high degree of accuracy. Also notice that the uncertainties begin small, but progressively increase throughout the run. What this means is as we progress, our life consumption estimates become increasingly uncertain. As discussed earlier, the large uncertainties can be largely attributed to the sparse training set where sampling can have a significant impact on the accuracy of the model estimates. The uncertainty bounds are expected to decrease as additional data is used to train the model.

5. CONCLUSION

Prognostics can be very valuable if it fits into current practices and is able to accurately predict failures. This paper has described a technique has been demonstrated to meet these two goals. The developed life consumption model was show to be able to estimate the life consumped of an individual drilling tool to within 4-12% with uncertainties of ±15-35%. The accuracy and uncertainty metrics for the developed model are expected to decrease as additional examples of device stress are used to train the model.

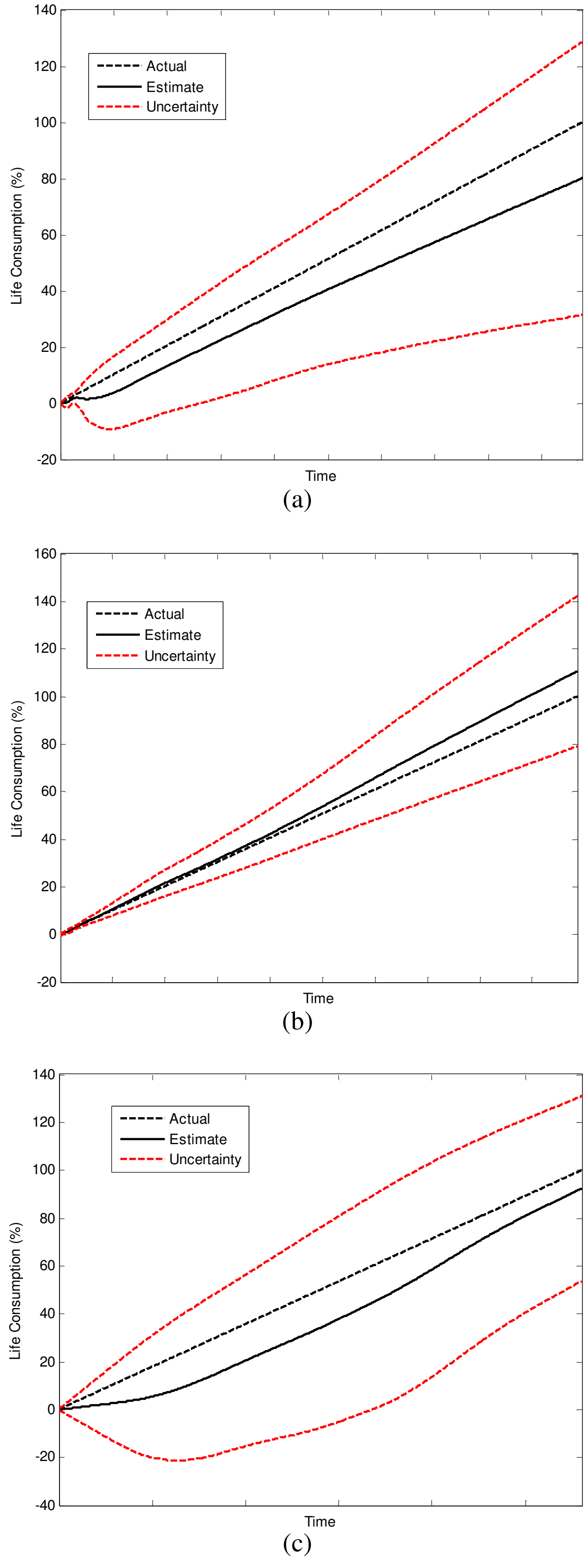

Figure 6. Modified PACE model estimates and uncertainties for the (a) first, (b) second, and (c) third test paths.

REFERENCES

(Brehme and Travix, 2008) J. Brehme and J. T. Travis (2008). Total BHA Reliability – An Improved Method to Measure Success, IADC/SPE 112644, *Proceedings of the IADC/SPE Drilling Conference*, Orlando, Florida.

(Efron and Tibshirani, 1994) B. Efron and R. J. Tibshirani (1994), *An Introduction to the Boostrap*, Monographs on Statistics and Applied Probability, Chapman & Hall/CRC.

(ExxonMobil, 2007) ExxonMobil (2007). ExxonMobil Announces Drilling of World-Record Well, Press Release: April 24, 2007.

(Fan and Gijbels, 1996) J. Fan and I. Gijbels (1996). *Local Polynomial Modeling and Its Applications*, Chapman & Hall/CRC, New York, NY.

(Garvey *et al.*, 2009) D. R. Garvey, J. Baumann, J. Lehr, and J. W. Hines (2009). Pattern Recognition Based Remaining Useful Life Estimation of Bottom Hole Assembly Tools, SPE/IADC 118769, *Proceedings of the 2009 SPE/IADC Drilling Conference and Exhibition*, Amsterdam, The Netherlands.

(Hines *et al.*, 2009) J. W. Hines, D. R. Garvey, R. Seibert, A. Usynin, and S. Arndt (2008), Technical Review of On-line Monitoring Techniques for Performance Assessment: Part II Theoretical Issues, NUREG/CR-6895, Vol. 2, United States Nuclear Regulatory Commission.

(Hines and Garvey, 2008) J. W. Hines and D. R. Garvey (2008). Nonparametric Model-Based Prognostics, *Proceedings of the Annual Reliability and Maintainability Symposium*, Las Vegas, NV.

(Tamhane and Dunlop, 2000) A. C. Tamhane and D. D. Dunlop (2000), *Statistics and Data Analysis from Elementary to Intermediate*, Prentice Hall, Upper Saddle River, NJ: 2000.

(Wand *et al.*, 2006) P. Wand, M. Bible, and I. Silvester (2006). Risk-Based Reliability Engineering Enables Improved Rotary-Steerable-System Performance and Defines New Industry Performance Metrics, IADC/SPE 98150, *Proceedings of the IADC/SPE Drilling Conference*, Miami, Florida.

Dustin R. Garvey is Diagnostics and Prognostics Team Lead at Baker Hughes. He received his B.S., M.S., and Ph.D. in Nuclear Engineering from the University of Tennessee in May 2005, December 2005, and May 2007 respectively. Dr. Garvey's research areas include developing data driven fault detection, diagnosis, and prognosis system for drilling systems. Dr. Garvey has authored 37 publications including one book entitled *Fuzzy Inference Based Monitoring, Diagnostics, and Prognostics*.

Martin John is an Electrical Engineer at Baker Hughes. He has a M.S. degree in electrical

engineering with specialization in process automation and microsystems technology. He currently works in the field of developing diagnosis and fault detection software for drilling systems.

Jörg Baumann is Director of Global Reliability and Fleet Management for Baker Hughes. He has a M.S. degree in mechanical engineering. He has worked for Baker Hughes in various engineering and management roles since 1988.