

Physics-based Remaining Useful Life Prediction for Aircraft Engine Bearing Prognosis

Nathan Bolander¹, Hai Qiu², Neil Eklund², Ed Hindle³, Taylor Rosenfeld³

¹ *Sentient Corporation, Idaho Falls, Idaho, 83404, USA*
nbolander@sentientscience.com

² *GE Global Research, Niskayuna, New York, 12309, USA*
qiu@research.ge.com
eklund@crd.ge.com

³ *GE Aviation, Cincinnati, Ohio, 45215, USA*
edmund.hindle@ae.ge.com
taylor.rosenfeld@ae.ge.com

ABSTRACT

Aircraft engine bearing prognosis not only requires early detection of a bearing defect, but also the ability to predict bearing health conditions for all operational scenarios. This paper summarizes a physics-based remaining useful life (RUL) prediction method developed in the DARPA Engine System Prognosis (ESP) program. This investigation focuses on a typical roller bearing fault (or defect) on the outer raceway. Spall detection is based on the fusion of vibration and online oil debris sensors. Spall size estimation is derived from the amount of bearing debris chips that passed through the Oil Debris Monitor sensor. Subscale propagation tests were performed to generate the response surface of the spall propagation rate under various operating speeds and loads. A particle filter based approach was used to track the spall propagation rate and update the prediction according to newly calculated diagnostics information. The bearing spall propagation model outputs a RUL distribution, which is calculated based on future operating conditions and the time the spall size crossing the failure threshold. The developed RUL prediction method was validated using full-scale bearing spall tests. The comparison of model prediction and measured ground truth demonstrated that the developed model was

able to predict the spall propagation rate accurately, and its prediction accuracy and confidence can be further improved by incorporating more diagnostics updates and/or increasing the confidence in the sensor data.*

1 INTRODUCTION

Engine bearing spalls are one of the leading causes of class-A mechanical failures leading to the loss of an aircraft (Wade, 2005). Bearing prognostics is the key to maximizing safety and asset availability while minimizing logistical costs, by allowing maintenance to be proactive rather than reactive (Marble and Morton, 2005). When a damaged or contaminated bearing spalls, metal particles from the bearing will eventually be detectable in the lubrication system. Today, bearing diagnostics is accomplished through the use of a Scanning Electron Microscope and Energy Dispersive X-Ray (SEM-EDX) examination of oil samples taken from the aircraft. The magnetic chip detector is examined for chips after every flight. If chips exist, they are removed and sent to the SEM-EDX machine to determine chemical composition and size. The approach is labor intensive, time consuming, and costly. Recent advances in sensor technology and computational intelligence have made real-time bearing prognosis feasible. However, bearing fault prognosis is

* This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

a very challenging subject, where two important questions need to be addressed. The first is to detect fault and assess its severity, i.e., where on the overall health curve the component or system resides (Brotherton, 2000). Once the current health condition is defined, the second question is to predict the change in component health as a function of RUL based on anticipated future missions.

Assuming a fault propagation model is available, it is possible to estimate a margin from the predicted failure threshold once the first question is addressed. However, it is hard to quantitatively diagnose the fault severity, especially at the early stage of fault. It is still very challenging to predict the future trend due to strong stochastic characteristics of the failure propagation process. And lastly, it is difficult to define a reasonable failure threshold, especially when limited historical failure data is available.

A large variety of RUL prediction algorithms have been proposed in the past research projects. Jardine et al.(2006) presented a review on various RUL estimation methods and defined the RUL as a conditional random variable of the time left before observing a failure given the current machine age and condition and past operation profile. It should be noted that the RUL is not only a function of past condition profiles, but up to future usage. As pointed out in (Jardine et al., 2006), RUL estimation methods fall into three main categories: statistical approaches (Wang, 2002, Banjevic and Jardine, 2005, Vlok et al., 2004, Phelps et al., 2001) artificial intelligent (AI) approaches (Gebrael et al., 2004, Zhang and Ganesan, 1997, Yam et al., 2001, Wang et al., 2004) and model-based approaches. Here, model-based approaches refer to specific physics-based fault propagation model, such as the spall propagation model to be discussed in this paper.

Statistical and AI approaches attempt to generate data-driven models to approximate the RUL distribution or the expectation value of RUL distribution. RUL estimation can be simplified as generating trending models, either from data-driven or in combination with degradation mechanism or empirical failure definition models, and use those models to forecast the degradation trend. A physics-based approach, such as the bearing spall propagation model, tracks damage accumulation and related accumulated plastic strain to cycle life. Unfortunately, they are often expensive to develop and are narrowly applicable. For instance, each physics-based model is specific to bearing material and geometry; any change to the bearing configuration requires the development of a new model. However, the advantage of this model is its capability to accurately factor in future operating conditions.

This paper summarizes a physics-based RUL prediction method for the aircraft engine bearing prognostics developed in the DARPA ESP Program. The model computes the spall growth trajectory and time to failure based on operating conditions, and uses diagnostic feedback to self-update, which reduces prediction uncertainty. Experimental data from the bearing test rig demonstrated that spall propagation rates can be predicted with higher confidence. This paper is organized as follows: Section 2 discusses the bearing diagnostics techniques, mainly focusing on how to detect a bearing spall and estimate the spall length using the fusion of vibration and oil debris data. Section 3 illustrates the steps of developing the spall propagation model, and how to use it to predict bearing RUL. Experimental bearing tests are presented in section 4. Finally, discussions and conclusions are provided in section 5.

2 BEARING DIAGNOSTICS

The architecture of the Bearing Prognosis Reasoner is shown in Figure 1. It consists of two major steps, diagnostics and prognostics. The objective of the diagnostic section is to determine the health of the bearing. If the model determines an unhealthy bearing exists by detecting the existence of spall signature and the spall size is determined using sensed vibration and oil debris data. The prognostic functions will be triggered by a positive diagnostic assessment. A key function in the prognostic module is the RUL calculation to determine the urgency of impending maintenance. This is accomplished by exercising the physics-based spall propagation model, which takes inputs of the initial spall size as well as future operating conditions and generates a series of possible spall propagation trends. The RUL distribution can therefore be approximated by computing when the propagation trend will pass a pre-defined failure threshold and trigger a maintenance action.

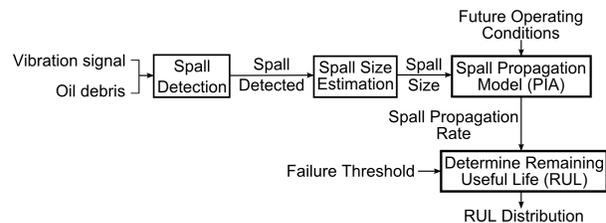


Figure 1: Bearing Prognosis Reasoner

2.1 Spall Detection

Vibration data monitoring is a widely used approach for spall detection. High frequency vibration features are known to be good indicators of incipient bearing defects. However, the performance is often influenced

by the load, operational speed, and background noise etc. On the other hand, an online oil debris sensor either captures metallic particles, or counts the particles, from which an estimate of total accumulated mass lost can be estimated. While metal particles in the oil system may indicate bearing spall, a complex aircraft engine lubrication system can trap a large fraction of particles. This limits the amount of mass detected by the sensor and delays or prevents the detection of a spall. Moreover, it is impossible to isolate which engine component is defective based solely on oil debris information, since other mechanical components may also shed metal particles under normal operating conditions. SEM-EDX is required at this point to identify the faulty component. Therefore, the fusion of vibration and oil debris information, capitalizing on the strengths of each approach, results in a sensitive and robust defect detection and isolation system.

A method called Synthesized Synchronous Sampling was used to convert the vibration data into an order domain and enhance the differential bearing damage signature. This technique, in combination with the conventional acceleration enveloping technique, allows the detection of inter-shaft bearing damage at a much earlier stage when compared to conventional enveloping analysis or spectrum analysis methods. For more information about this method please refer to (Luo and Qiu, 2009).

A fuzzy logic based sensor fusion scheme was developed to integrate the vibration feature with oil debris data for spall detection. Fuzzy membership functions and fuzzy rules were derived from the experimental data. The output of the spall detection fusion module is a detection flag. A value of 1 indicates a spall is detected. Once a spall is detected, the next step commences a quantitative estimation of initial spall size.

2.2 Quantitative Bearing Diagnostics

The initial spall size is the estimated spall size at the moment when the spall is initially detected. A spall size estimate is one of the initial conditions required to run the spall propagation model. The spall size estimation algorithm relies on the oil debris sensor that provides a monotonic signal related to spall length. Assuming a faulty bearing component, it is possible to estimate spall length from total chips counted in the scavenge line. The magnitude of vibration features, however, are often less clearly related to damage magnitude – it is not unusual to see the magnitude of a frequency domain vibration feature decrease as the fault develops due to the stochastic nature of fault development and the occasional shift in energy at different stages of failure.

Multiple rig tests were conducted to derive the relationship between the quantity of detected oil chips and the actual spall size on the bearing raceway. Figure 2 shows the test data from multiple rig tests and the derived linear spall size estimation model as well as the 95% confidence intervals and 95% predictive intervals.

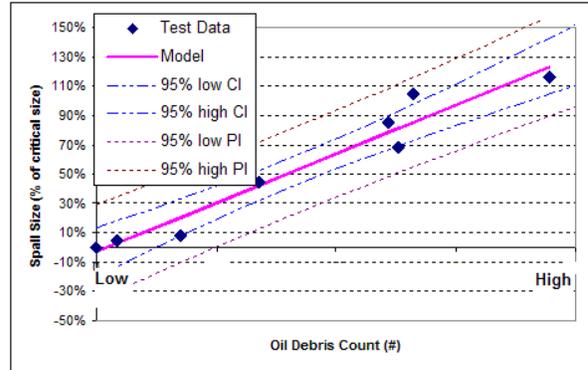


Figure 2: Spall size estimation based on oil debris data

3 Bearing Prognostics Integration Architecture

3.1 Bearing Failure Threshold - Critical Spall Length

Figure 3 depicts the general process for spall initiation and propagation. Critical spall length is defined as length of one rolling element spacing (actual length varies depending on the bearing).

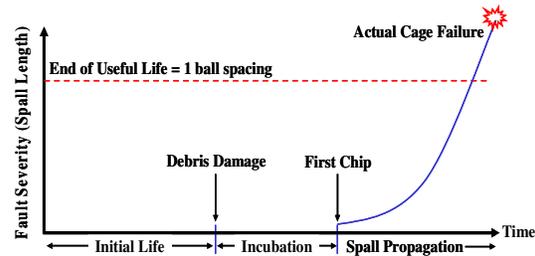


Figure 3: Bearing spall propagation process

Destructive failure occurs when the cage fails. In low speed bearings, a spall can cover the entire race without catastrophic failure. However, for high-speed bearings, the bearing reaches end of life when the spall length is greater than the circumferential ball or roller spacing. Stress on cage crossbars and rails increases dramatically when spall length is greater than ball spacing and a cage failure usually happens soon afterwards.

3.2 Physics of Bearing Spall Propagation

Bearing spall propagation is a complex phenomenon describing the growth of an existing damaged region on

a bearing race or roller due to the quasi-continuous liberation of material in the form of chips/particles during operation. The rate of damage accumulation due to a given set of operating conditions (speed/load) is dependent upon the properties of the particular material under consideration. In this study, subscale seeded-fault spall propagation testing was utilized to investigate the spall propagation phenomena and determine the rolling contact fatigue/impact damage accumulation behavior of the bearing material.

Figure 4 shows the test rig developed as part of this study for subscale spall propagation testing of cylindrical roller bearings. Obtaining experimental data of sufficient density is quite labor intensive. The key feature of this test rig is the ease with which the test bearing can be removed for inspection.



Figure 4: Subscale spall propagation test rig

Inspection photographs, along with periodic samples of vibration and temperature data, are stored in a proprietary database.

Figure 6 presents an example case of spall propagation from a seeded fault. There are three distinct phases of spall growth. At first, the spall grows slowly (Figure 6, frames 1-3), undergoing an 'incubation' period prior to downstream propagation. During the incubation phase, growth is primarily outwards rather than downstream.

While there is a solid qualitative understanding of spall behavior during the incubation phase, quantitative spall size estimates during this phase present some challenges. Growth during this phase is governed primarily by two factors: a) the degree of damage imparted at initiation (i.e. the magnitude of plastic deformation and residual stresses at the indentation) and b) the development of the initial subsurface crack network. Both of these factors are subject to some degree of uncertainty, and are not considered in the current analysis.

Once the spall has propagated across the race, it transitions to the 'propagation' phase (Figure 6, frames 4-9). This phase of spall growth is typically regular

and well behaved enough to enable predictive modeling, and is therefore the focus of the current modeling effort. Growth during this phase is driven by the impact and reloading of the roller as it reaches the trailing edge of the spall. This leads to an accumulation of plastic strain in the material, ultimately resulting in crack propagation and chip liberation.

After the cumulative spall length has surpassed a certain threshold, the rate of propagation accelerates significantly (Figure 6, frames 10-12). The transition to this 'accelerated growth' phase is typically defined as the failure threshold. In high-speed applications this transition point is associated with a spall length corresponding to one roller spacing, after which cage failure occurs, leading to catastrophic failure of the bearing.

Figure 7 illustrates the downstream spall growth during the 'propagation', and early 'accelerated growth' phases observed during several tests.

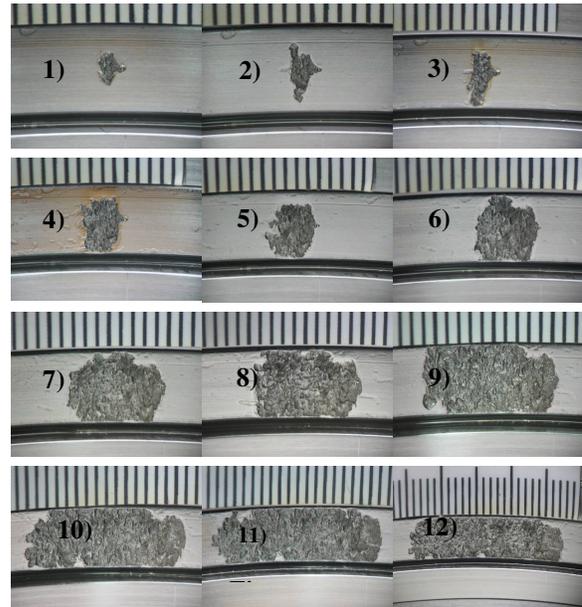


Figure 6: Spall propagation for a cylindrical roller bearing under radial loading. Rolling direction is right-to-left

Each of these tests were run under different (constant) operating conditions (constant load & speed), with data set TS03 corresponding to the results presented in Figure 6). For comparison purposes the time axis in Figure 7 has been normalized with respect to the time at which the spall reached approximately 20mm in length. An interesting knee in the curve can be observed in Figure 7, roughly corresponding with the spacing between rollers. Spall growth prior to this point (the 'propagation' phase described previously) is typically very regular, lending itself well to predictive modeling. The 'accelerated growth' phase beyond this

knee is similarly well behaved for these bearings, however operation in this region carries an increased risk of cage failure.

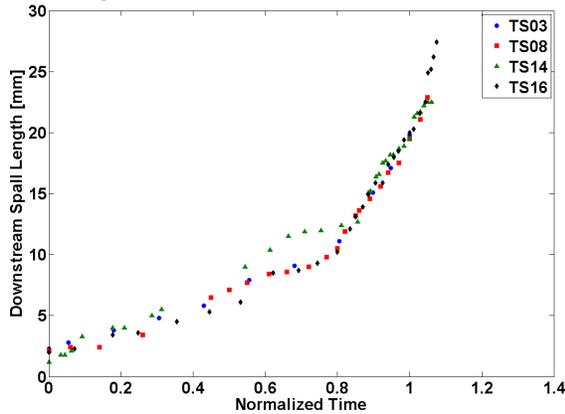


Figure 7: Spall growth during propagation phase

3.3 CABPro Model

There are two key elements required to model spall propagation: determination of dynamic loads and stresses occurring in the material as a rolling element passes over the spall, and development of a method relating this local stress field to damage accrued in the material. Sentient Corporation has developed a physics-based model to predict the rate at which bearing spall damage will progress under a given set of operating conditions. CABPro tracks the material state using principles of continuum damage mechanics, which relates to localized stress and strain to microstructural degradation and eventual failure in a widely applicable way. The purpose of the subscale testing described in the previous section was to characterize the behavior of the bearing material under rolling contact fatigue/impact. In a damage mechanics model (such as the one in CABPro) the behavior of the material is described by one or more *material parameters*. In this case, the material parameters are embedded in the rate of spall propagation that is measured periodically during the subscale tests via teardown and inspection. Through accumulation of spall propagation data over a range of operating conditions, the parameters characterizing the RCF/impact damage behavior of the bearing material can be extracted and applied in the continuum damage mechanics approach for the full scale bearing. Further details of the CABPro model can be found in (Marble and Morton, 2006)

The contact-level conditions are based on the overall loads, speeds, lubrication, etc. applied to the bearing. The accumulation of damage during a stress cycle is related to the existing damage and to the applied stress via the incremental plastic strain energy accumulated. A custom damage accumulation program imports stress and strain data from finite element analysis (FEA) and applies damage mechanics to

calculate the spall propagation rate for a particular geometry and load/speed combination.

Figure 8 presents a 3D finite element model of a segment of a cylindrical roller bearing with a simulated spall. Cyclic boundary conditions are applied at each end of the segment, with the assumption that the interaction between the roller and spall is sufficiently localized so as not to propagate to the neighboring segment. A symmetrical boundary condition is applied along the rotational axis to exploit the plane of symmetry in the bearing, thus requiring that only half of the total geometry be modeled. The materials of the inner race and cage are defined as rigid, with prescribed motion given by the kinematics of the bearing. These two components act to drive the roller through the spalled region and reloading zone. The local stresses generated during the impact/reloading event at the trailing edge of the spall are the principle driving mechanism behind spall propagation. The stresses from the FE analysis are imported into CABPro and used to calculate the rate of damage accumulation in the material near the trailing edge. This analysis is repeated periodically to update the stress/strain fields as the crack network grows and the surface material erodes.

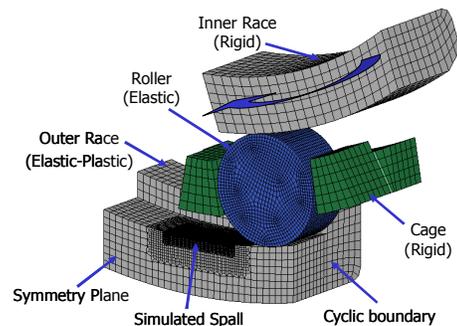


Figure 8: FEA modeling of roller/spall impact

Figure 9 depicts the formation of a chip at the trailing edge of a simulated spall. The removal of material in discrete amounts explains the linearity of the spall propagation curve – the material removal process is quasi-continuous. Damage due to the impact/reloading event is confined to a small region surrounding the trailing edge of the existing spall. Essentially, spall propagation is a self-resetting fatigue process; a new region of material is exposed to the impact/reloading event as the previous damaged material is removed. Thus, while the damage accumulation at the trailing edge is a decidedly non-linear process, the spall propagation process, on average, is linear. Once the parameters in the damage equation have been calibrated such that they adequately represent the rolling contact fatigue/impact behavior of the bearing material, the CABPro model can be used to explore RUL under various mission load spectrums.

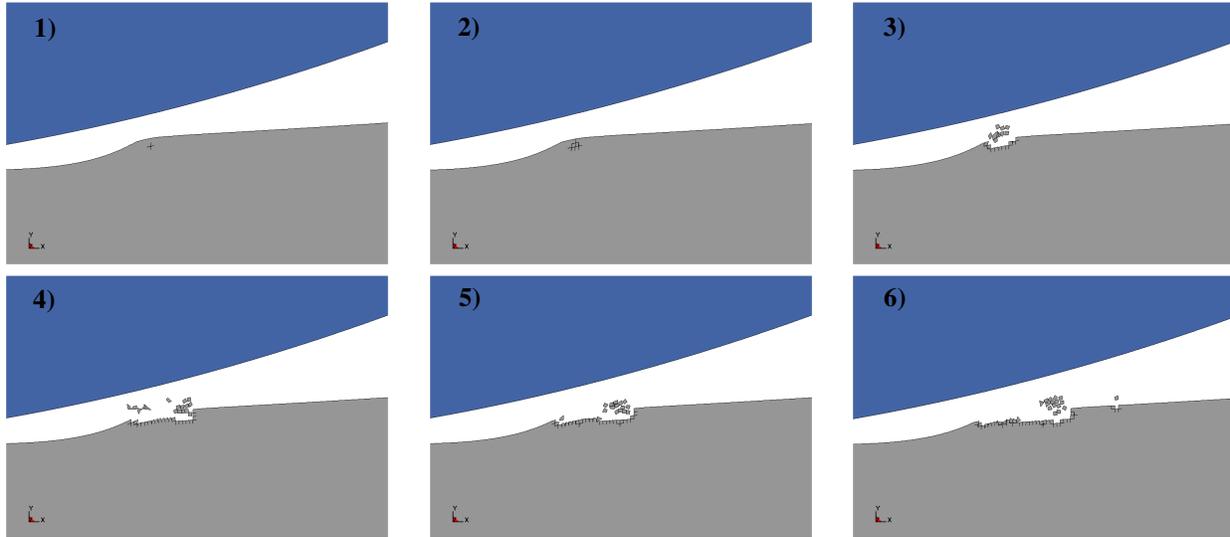


Figure 9: Snapshots from FE runs showing evolution of damage at the trailing edge of an existing spall

3.4 Spall Rate Response Surface

The full high-fidelity CABPro model is very computationally intensive due to the iterative use of finite element analysis to determine the local stress fields. However, due to the linear behavior of the spall growth during the ‘propagation’ phase (as illustrated in Figure 6), the behavior of the high-fidelity CABPro model can be captured by a response surface describing the rate of spall propagation as a function of load and speed. This reduced order model (ROM) can then be used within the online model updating procedure described in the next section.

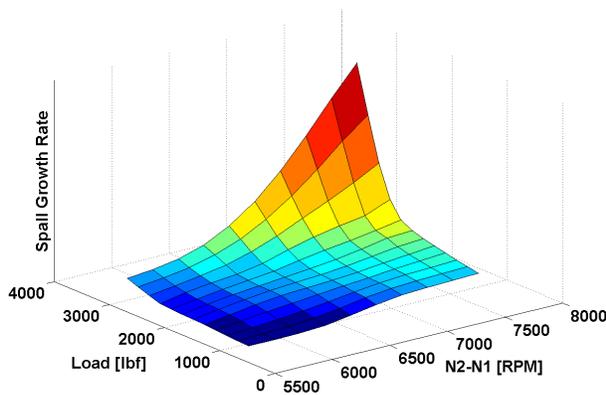


Figure 10: Spall rate response surface

The first step in developing the spall rate response surface is to define the operational envelope for the bearing. This envelope should include all anticipated potential load/speed combinations encountered by the bearing during operation. The full CABPro model exercised at a sufficient number of points within this

region to map out the spall propagation rate response surface. Figure 10 present the response surface developed for this study.

The response surface presented in Figure 10 exhibits the expected behavior, i.e. spall growth rate increases with operating condition severity, with a slightly greater dependence upon load than speed

3.5 Online Model Updating

Model updating refers to the process of utilizing diagnostic data as a source of additional knowledge in order to reduce uncertainty in the RUL prediction. Proper model updating approaches view the model as a general description of fault progression characteristics and the sensor based diagnostics as a noisy indication of current state. An improved estimate of state can be obtained by combining the sensor-based state estimates with a fault or damage progression (prognosis) model. Sentient Corporation’s Prognostic Integration Architecture (PIA) is a stochastic framework and set of general-purpose algorithms for fusion of diagnostic state indications with damage progression models. It provides an automated prediction of current state and remaining life with accurate and optimal uncertainty bounds. The PIA is a mature, generalized architecture applicable to a wide range of diagnostics and prognostic models at the component level. In this section, a description of the methodologies employed by the PIA will be provided, followed by a discussion of integration with diagnostic and damage progression models for an example dataset.

The PIA is based on a particle filter approach with Bayesian updating. Particle filters are most commonly used to directly estimate the observable state of interest, which for this application is the spall severity. In the

PIA, Sequential Monte Carlo methods have been developed to indirectly estimate state by employing them in a parameter identification mode. The parameter(s) to be identified are initially unknown constants that describe the *differences* in damage propagation behavior between individual components. The objective of the model updating scheme is to reduce uncertainty in both current state estimates and forward predictions by *learning* the characteristics of an individual component as it degrades. This model updating scheme is flexible, powerful, and applicable to a large class of problems in health management and prognostics.

Figure 11 provides a conceptual illustration of the parameter identification process. Stochastic parameters are used to represent the difference between the “average” component and a particular component. Based on the damage progression model and the uncertainty of the parameters, a group of particle values is sampled from the parameter distributions; each sampled value is used to generate a candidate damage trajectory (particle model).

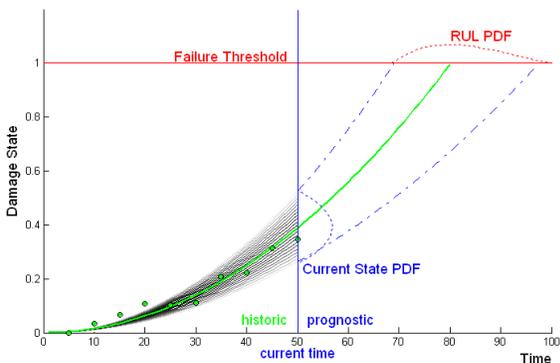


Figure 11: Conceptual illustration of model updating via Bayesian fusion

The black lines in Figure 11 represent the range of possible damage trajectories, the variance in the trajectories being reflective of the sampled particle parameter values. By utilizing a source of additional information (diagnostic data, green points), the values of the particle parameters (and the resulting damage trajectory) that best represent the particular unit under test can be ascertained.

The challenge is to determine which, among a family of trajectories as defined by the damage rate parameter, best represents the particular bearing under consideration. This is accomplished by incorporating additional information obtained through the incoming diagnostic data. A Bayesian updating procedure is used to weight the particle trajectories based on how well they fit the incoming and past diagnostic data. These fitness values, or weights, are applied to a (Gaussian)

kernel function for each particle, which are then combined in a Gaussian mixture density to provide a probability distribution for the current state of damage. This current state distribution is then propagated into the future past the failure threshold to determine the RUL distribution.

4 EXPERIMENT

4.1 Rig Setup and Data

An existing engine rig located in GE’s bearing lab was modified to suit the bearing spall test experiment and installation requirements. The full scale bearing test rig has the capability for periodic partial teardowns to inspect photograph and measure spalls generated in the outer race. Multiple rig tests had been conducted. One test result was chosen for discussion here. To create a realistic spall initialization condition, the test bearing started with one indent on the outer race. A template centered the indent in the middle of the roller path. Four inspections including a final teardown were performed at total run time (TRT) of 1.64, 38.6, 48.8, and 57.1 hours respectively (see Figure 12). A spall with the size estimated to be 0.037 square inches was detected at TRT 38.6 hour.

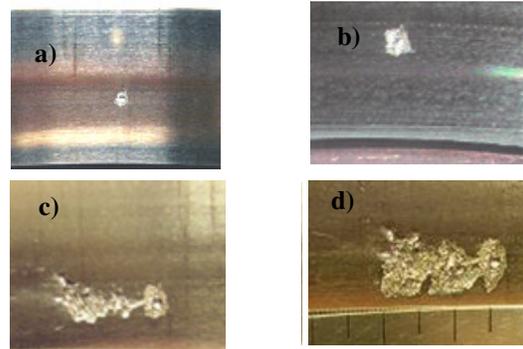


Figure 12: Inspection at TRT (a) 1.64 hrs (b) 38.6 hrs (c) 48.8 hrs (d) 57.1 hrs

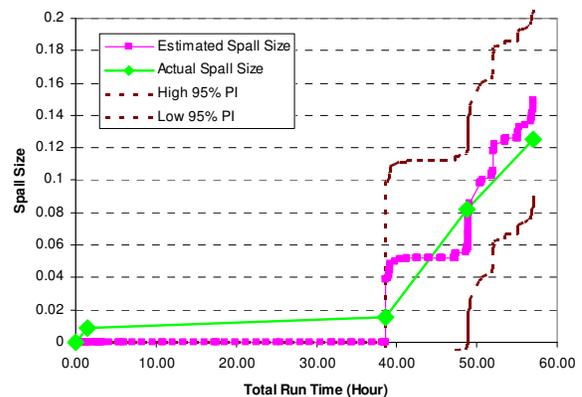


Figure 13: Spall estimation vs. actual tear down measurements

Subscale seeded-fault spall propagation testing was utilized to investigate the spall propagation phenomenon and determine the rolling contact fatigue/impact damage accumulation behavior of the bearing material. A custom damage accumulation program imports stress and strain data from FEA and applies damage mechanics to calculate the spall propagation rate for a particular geometry and load/speed combination. Based on the subscale spall propagation tests and FEA analysis, a response surface was then developed to describe the rate of spall propagation as a function of load and speed.

Another unique feature of the developed physics-based RUL prediction method is the model updating function, which refers to the process of utilizing diagnostic data as a source of additional knowledge in order to reduce uncertainty in the RUL prediction. The RUL prediction method is based on a particle filter approach with Bayesian updating. The prognostic model is first initialized using a-priori, expert knowledge. A Bayesian updating procedure is used to weight the particle trajectories based on how well they fit the incoming and past diagnostic data. These fitness values, or weights, are applied to a (Gaussian) kernel function for each particle, which are then combined in a Gaussian mixture density to provide a probability distribution for the current state of damage. This current state distribution is then propagated into the future – beyond the failure threshold – to determine the RUL distribution.

The developed RUL prediction method was validated by a full-scale bearing test. Comparison of model prediction and measured ground truth demonstrated that the developed model was able to predict the spall propagation rate accurately, and its prediction accuracy and confidence can be further improved by incorporating more diagnostics updates and/or increasing the confidence in the diagnostic data.

ACKNOWLEDGMENT

The experimental data used in this paper was obtained from the work supported by the Defense Advanced Research Projects Agency, Defense Sciences Office (DSO), Engine System Prognosis, issued by DARPA/CMO under Contract Number: HR0011-04-C-0002. Authors would also like to express gratitude for the insightful suggestions and guidance given by the DARPA program officer and the Red Team.

REFERENCES

D. Banjevic and A.K. Jardine. (2005). Calculation of reliability function and remaining useful life for a Markov failure time process. *IMA Journal of Management Mathematics*.

- T. Brotherton. (2000). Prognosis of Faults in Gas Turbine Engines, in *Proceedings of IEEE Aerospace Conference*, vol. 6, pp. 163-171.
- N. Gebraeel, M. Layley, R. Liu, and V. Parmeshwaran. (2004). Residual Life Prediction from Vibration-based Degredation Signals: A Neural Network Approach, *IEEE Transactions on Industrial Electronics*, vol. 51(3), pp. 694-700.
- A.K. Jardine, D. Lin, and D. Banjevic. (2006), A review on machinery diagnostics and prognostics implementing condition based maintenance, *Mechanical Systems and Signal Processing*, vol. 20(7), pp. 1483-1510.
- H. Luo, H. Qiu et al., (2009) Synthesized Synchronous Sampling Technique for Bearing Damage Detection, in *Proceedings of the ASME 2009 IDETC/CIE*, Aug. 30 – Sep. 2, 2009, San Diego, California, USA.
- S. Marble and B. Morton (2005), Predicting the Remaining Life of Propulsion System Bearings, in *Proceedings of IEEE Aerospace Conference*, Big Sky, MO.
- E. Phelps, P. Willett, and T. Kirubarajan. (2001). A statistical approach to prognostics. in *Component and Systems Diagnostics, Prognosis and Health Management*, vol. 4389, Bellingham pp. 23-34.
- P.J. Vlok, M. Wnek, and M. Zygmunt. (2004). Utilizing statistical residual life estimates of bearings to quantify the influence of preventative maintenance actions. *Mechanical Systems and Signal Processing*, vol. 18, pp. 833-847.
- R.A. Wade. (2005), A Need-focused Approach to Air Force Engine Health Management Research, in *Proceedings of IEEE Aerospace Conference*, Big Sky, MO.
- W. Wang. (2002). A model to predict the residual life of rolling element bearings given monitored condition information to date, *IMA Journal of Management Mathematics*, vol. 13, pp. 3-16.
- W.Q. Wang, M.F. Golnaraghi, and F. Ismail. (2004). Prognosis of machine health condition using neuro-fuzzy systems, *Mechanical Systems and Signal Processing*, vol. 18, pp. 813-831
- R.M. Yam, P.W. Tse, L. Li, and P. Tu. (2001). Intelligent predictive decision support system for condition based maintenance, *International Journal of Advanced Manufacturing Technology*, pp. 383-391.
- S. Zhang, and R. Ganesan. (1997). Multivariable trend analysis using neural networks for intelligent diagnostics of rotating machinery, *ASME Journal of Engineering for Gas Turbines and Power*, vol. 119, pp. 378-384.



Nathan Bolander is a Research Engineer at Sentient Corporation and serves as the lead developer for Sentient's Component Life Prediction (CLP) software. He holds an M.S. (2002) and Ph.D. (2007) in Mechanical Engineering from Purdue University, with an emphasis in tribology. His research interests include physics-based modeling of surface interaction, damage progression, and lubrication in rolling/sliding contacts such as those found in bearings and gears.



Hai Qiu is a lead research scientist in the GE Global Research at Niskayuna, New York. Prior to joining GE in 2005, he was a Research Assistant Professor in the Department of Mechanical, Industrial and Nuclear Engineering of the University of Cincinnati and served as the Lead Researcher of the NSF Industrial/University Cooperative Research Center for Intelligent Maintenance Systems (IMS). He obtained his Bachelors and PhD degrees in mechanical engineering from the Xi'an Jiaotong University in 1995 and 1999, respectively. He has conducted a wide variety of research projects in the fields of prognostics and intelligence maintenance systems, funded by the NSF and industry. His current research areas include intelligent diagnostics and prognostics, advanced signal processing, and applied artificial intelligence.



Neil Eklund received a B.S. in 1991, two M.S. degrees in 1998, and a Ph. D. in 2002, all at the Rensselaer Polytechnic Institute. Dr. Eklund was a research scientist at the Lighting Research Center from 1993 to 1999. He was in the network planning department of PSINet from 1999 to 2002, before joining General Electric Global Research in Niskayuna, NY in 2002. He has worked on a wide variety of research projects, including early detection of cataract using intraocular photoluminescence, multiobjective bond portfolio optimization, and on-wing fault detection and accommodation in gas turbine aircraft engines. His current research interests involve

applying artificial intelligence to create robust solutions to real-world problems for the finance, aviation, and oil and gas industries. Dr. Eklund is also an adjunct professor in the Engineering/CS department at Union Graduate College in Schenectady, NY, since 2005 where he teaches classes in Computational Intelligence and Machine Learning.



Ed Hindle is the Sustainment and Prognosis Health Manager at General Electric Aviation Advanced Technology and Preliminary Design group. He manages a team of engineers responsible for developing advanced prognostic technologies to better understand and accurately manage engine system health for commercial and military applications. He has a B.S. in Mechanical Engineering, a Master of Business Administration from the University of Miami and a Master of Divinity from Denver Seminary. He has over 27 years of aerospace experience in materials evaluation and research, military life management, advanced technology program management for the department of defense and engine systems integration. Currently he is the program manager for the DARPA Engine System Prognosis Program, Propulsion Safety & Affordable Readiness Initiative, Engine Rotor Life Extension Program, and Agile Combat Support programs.

Taylor Rosenfeld Taylor a graduate of California State University, Fresno with a B.S. Degree in Mechanical Engineering. He began his career in 1979 on the Engineering Development Program at GE Aircraft Engines. He has held various positions of increasing responsibility in aerodynamic-thermodynamic analysis and systems engineering on many GEAE products. In 1993, he was promoted to Manager of Performance Engineering responsible for preliminary design, development, certification, and field support for the CF6 product line. In March 2000, Taylor was appointed Six Sigma Master Black Belt for Military Inlets & Exhaust Systems, responsible for integrating quality processes, leading the design for six sigma initiative, and mentoring Black Belts. He was promoted to manager of the Intelligent Engine program in November 2001.