

USING SIMULATION, DATA MINING, AND KNOWLEDGE DISCOVERY TECHNIQUES FOR OPTIMIZED AIRCRAFT ENGINE FLEET MANAGEMENT

Michael K. Painter
Madhav Erraguntla
Gary L. Hogg, Jr.

Knowledge Based Systems, Inc.
1408 University Drive East
College Station, TX 77840, U.S.A.

Brian Beachkofski

Air Force Research Laboratory (AFRL/PRTS)
Wright-Patterson AFB, OH 45433, U.S.A.

ABSTRACT

This paper presents an innovative methodology that combines simulation, data mining, and knowledge-based techniques to determine the near- and long-term impacts of candidate aircraft engine maintenance decisions, particularly in terms of life-cycle cost (LCC) and operational availability. Simulation output is subjected to data mining analysis to understand system behavior in terms of subsystem interactions and the factors influencing life-cycle metrics. The insights obtained through this exercise are then encapsulated as policies and guidelines supporting better life-cycle asset ownership decision-making.

1 INTRODUCTION

The pressures of sustained high tempo military operations, spiraling maintenance costs, and concerns over rapidly aging aircraft assets have generated renewed interest in improved processes and technologies that can dramatically improve both affordability and readiness. A largely untapped source of significant gains in these areas is a set of carefully designed fleet maintenance policies and practices guided by a clear understanding of the consequences of today's decisions on near- and long-term performance. Simulation plays a potentially significant role in this discovery process.

In traditional simulation, system behavior is usually analyzed and interpreted in terms of statistical metrics (e.g., mean, variance, confidence intervals for throughput, queuing time, utilization) (Benjamin 1995, Erraguntla 1994). Simulation has also been used in conjunction with optimization to seek optimal parameter values or to do tradeoff analyses (Goldberg 1989). A more expansive framework is needed, however, to fully understand the dynamics of this system's behavior and the factors influencing overall performance. A key innovation of the approach presented in this paper is the marriage of data mining and

knowledge discovery techniques with simulation to develop the necessary insights and project the future conditions that will emerge as a consequence of today's decisions (Painter 2005). The methodology also uniquely supports explicit consideration of the context in which maintenance decisions are or will be made by directly leveraging asset status and maintenance history data. Together, this framework provides for significantly enhanced, measurable improvements.

2 AFFORDABLE READINESS CHALLENGE

A growing concern among Department of Defense (DoD) decision-makers is the accelerating cost of operations and sustainment (O&S). A significant portion of this cost is directly linked to the 47,630 turbine engines that power the 24,350 military aircraft in inventory. Sustainment costs for these assets exceed \$6.66B every year, excluding fuel expenses (Beachkofski 2006). And these costs are growing. Over the last ten years, the Navy experienced a 13.3% annual growth in cost per engine flight hour (CPEFH) (Cifone 2006). The Air Force reports that its maintenance costs are growing at twice the rate of inflation.

As O&S costs increase, it becomes too costly to invest in new systems. Consequently, aging weapon systems, many of which have been in service well beyond their design life, are pressed into continued service. As these systems continue to age, they begin to experience new and more frequent failures. Recapitalization, engineering modification, and life extension programs increase O&S costs while consuming proportionally higher percentages of that budget. Meanwhile, fleet modernization plans have been severely curtailed or put on hold simply to cover the costs of sustaining the assets that are currently in inventory.

The impact of these trends on mission capability is an even greater concern. The Navy, for example, lost six aircraft and seven aviators over a span of two months in early 2006, precipitating aviation standdown orders by the

Commander of Naval Air Forces in San Diego for the first time in four years. (CNAF 2006). By this time, the Navy had experienced four Class A flight mishaps attributed to engine-related problems in the first half of the fiscal year, exceeded their annual average (Cifone 2006). A Class A mishap is defined as one involving an aircraft with an intent to fly, which resulted in damages totaling \$1 million or more, a destroyed aircraft, a fatality, or a permanent total disability.

In view of these trends, joint service goals have been set to reduce propulsion-related Class A mishaps by 75%, reduce maintenance costs by \$420M (10%), and double the average time on wing for aircraft engines.

There is clearly a tension here among competing objectives. Many view the Reliability Centered Maintenance (RCM) concept as a possible way to reverse the trend of O&S costs spiraling out of control while simultaneously ensuring high levels of readiness and safety.

3 RELIABILITY CENTERED MAINTENANCE OVERVIEW

Current engine maintenance practices across the services involve frequent inspections, parts replacements, and rework in an attempt to maintain high standards of system reliability. Frequent maintenance correspondingly requires a large inventory of spare engines and engine components. Some studies (Nowlan 1978) indicate that current maintenance practices are not cost-effective, and may even reduce reliability and availability.

RCM provides a method for rational decision-making that considers the overall effects of possible maintenance actions — or the risk of not performing those maintenance actions — on aggregate level metrics like engine availability, performance, and life-cycle cost (LCC).

Recent efforts within the DoD to apply RCM seek to leverage engine maintenance and reliability history and cost data to drive engine maintenance decision-making such that CPEFH is minimized. The goal is to simultaneously reduce costs (i.e., improve affordability) and increase time on wing (i.e., readiness). Estimating the CPEFH implications of a given maintenance decision requires developing projections of expected time on wing (ETOW) and LCC. Maintenance history data is used to fit models yielding ETOW projections. The specific process by which ETOW estimates are generated is beyond the scope of this paper. Suffice it to say that these estimates are generated using Weibull distributions derived from field data reflecting aircraft engine maintenance history and reliability characteristics. Today, RCM decision makers are only able to estimate the immediate costs of a maintenance decision although they recognize that there are long-term cost consequences for those decisions. The following section describes the unique challenges and methodology developed to estimate LCC.

4 LIFE-CYCLE COST ESTIMATION METHODOLOGY

Every maintenance decision, including the decision to defer scheduled or unscheduled maintenance effort, has both immediate and long-term cost implications. These are linked decisions with probabilistic outcomes whose cost implications can only be understood by projecting the combined effects of several factors. Accurate estimates of the LCC implications of a given maintenance decision are needed to achieve an optimal balance among competing enterprise objectives. Particularly for aircraft engines — whose service life can span decades — maintenance decision-making that is driven to minimize immediate costs alone will tend to increase LCC and reduce operational availability.

Future costs based on historical data are likely to be poor estimators for several reasons. For example, sources of historical costs provide accounting data that is known to reflect a skewed picture of true costs. Furthermore, since engines are essentially perpetual life assets, there is little or no data reflecting assets that have completed a full life-cycle maintenance history and cost accumulation. Finally, historical patterns of maintenance do not necessarily depict what will be experienced in the future, particularly if planned usage patterns and the environment of use changes.

An important contribution of this research was the development of a methodology for LCC estimation that addresses these issues by coupling discrete event simulation and data mining to develop and maintain meaningful cost estimation parameters. The process involves simulating the fleet of engines for some planned set of operational scenarios to collect maintenance decisions and their cost implications over their entire service life. Multiple simulation runs are executed to account for the stochastic nature of maintenance events and decisions. This simulation-based maintenance and cost history data is then mined to discover the parametric relations that best characterize the LCC implications of each maintenance decision. The cost parameters and/or models developed through mining this data are then refined and eventually released for use in RCM decision-making. Maintaining these parameters requires periodic refinements motivated through changes in environment, maintenance practices, the underlying reliability of engine components, etc.

The logic of the simulator emulates the dynamics of engine fleet operation and maintenance over time using the following process:

1. Determine decision points (i.e., the times at which maintenance decisions must be made).
2. Determine decision options at the decision point.
3. Select decision option.

4. Determine costs incurred with the selected decision.
5. Iterate, returning to step 1 until a stopping condition is met (e.g., end of time horizon).

In the first step, the simulator determines the points during the life cycle of the asset at which the next set of maintenance decisions will be made. Some decision points are determined based on calendar-driven requirements (e.g., aircraft Programmed Depot Maintenance), while others are usage-driven as determined by hard life limits, engine reliability characteristics, optempo, and/or the operational environment (Figure 1).




Figure 1: Decision Point Timeline

Once the decision points are identified, the simulator must determine the range of possible maintenance actions.

Once a decision option is selected, the costs incurred as a consequence of that decision are determined. These costs may include labor costs, material costs, transportation costs, foregone revenues due to loss of availability, etc. This process is repeated a sufficient number of times to generate maintenance history and cost statistics that can be studied using data mining.

Simulation affords the opportunity to do what cannot be done even with the best possible maintenance history data collection system. It allows one to produce data reflecting operation of the engine fleet throughout its entire service life for any set of conditions. Simulation also allows one to determine the combined effects of variability and competing constraints on LCC performance.

Once this data is generated, data mining techniques can be used to identify key patterns and parametric relationships. These can then be used to define models from which LCC can be determined parametrically; thereby yielding decision support that promotes both accuracy and ease of use.

5 ENABLING TECHNOLOGIES

The technology needed to generate these LCC estimates includes a discrete event simulation-based cost projector, or Cost Projection Simulator (CPS) and a suite of data min-

ing tools used to develop and maintain LCC estimation parameters and/or models.

The CPS emulates the time varying cost behavior that will be manifested for an engine to support planned missions in the context of some set of maintenance policies. As the engines are “flown” with the reliability characteristics of the engines generating simulated maintenance events. The cost consequences of simulated maintenance events are recorded along the way until the specified stopping condition is met (e.g., engine reaches its service life limit). Differences in LCC among alternative maintenance decisions are determined by simulating alternative maintenance options and examining how they compare. Simulation runs are repeated multiple times for each decision option to establish expected cost and confidence intervals.

The CPS model encompasses all of the reliability modeling aspects needed for reasonable estimation of the costs associated with an engine life cycle. In this role, the CPS is designed to generate estimations of LCC based on the key factors that drive their structure. In essence, the model “flies” the fleet of engine assets to subject them to the accumulation of age, damage, and maintenance/repair actions that keep them operational. This goes beyond just simulating the accumulation of run time on an engine and the occurrence of scheduled maintenance events. The model also uses characteristic life functions to model the potential failures of the modules and the engine during operation. The underlying Weibull distributional forms are an integral piece of the model structure, and the parameters for each module are carried with it throughout its life in the model. Also embedded in the model structure are the hard-life limits for each component, as well as a representation of the maintenance policies and actions that these limits impose. The model also incorporates a number of other factors that are relevant to determining the LCC of an asset (see Figure 2). These parameters are not static. As the asset ages through use, these parameters are updated to reflect imposed changes. The methods for adapting these parameters include consideration of operational tempo, mission profiles, and operating environment.



Figure 2: Simulation-Based Cost Projector

6 SIMULATION MODEL AND PRELIMINARY DATA MINING FINDINGS

6.1 Simulation Model

The CPS was designed as a discrete-event simulation model constructed in Arena. This model had two primary purposes. The first purpose was to support long-range cost estimates. The second purpose was to generate data that could be subjected to data mining analysis to determine underlying cost relationships and structures.

The CPS model is structured to simulate the flying, removals, and maintenance actions to which an engine is subjected over its life cycle. Of primary importance in simulation are the embedded reliability characteristics of the engine components. For this model, the representative Weibull characteristics for each module are modeled to generate the appropriate failure modes and to determine the maintenance requirements that the asset will be subjected to over time. Statistical sampling is used to determine when a failure will be experienced, and, if so, what module(s) is involved, and the nature of the failure. If the engine completes a mission successfully, the components are aged according to the mission profile that was flown, and the parameters for the module reliability are updated. If the engine experiences a failure, the nature of the failure determines the repair and replacement options available together with which action is taken. The implementation of this model in Arena is depicted in Figure 3.

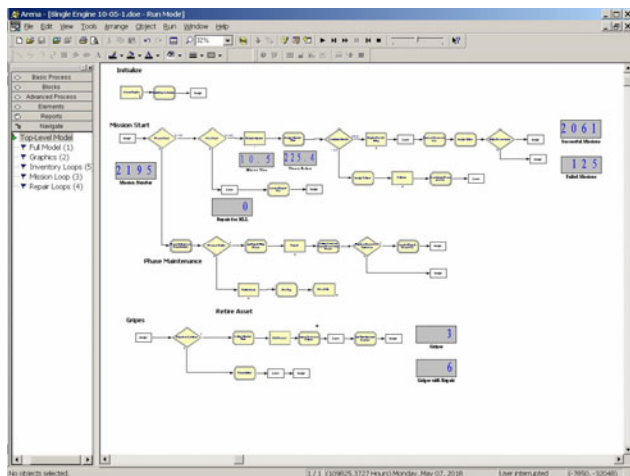


Figure 3. Cost Projection Simulator

The data generated by the simulator (part of which is shown in Figure 4) allows for detailed examination of LCCs and how they are accumulated over time. This provides a history of all maintenance actions and their associated costs. From this data, the long-term cost consequences of particular maintenance actions can be extracted. In addition to individual costs, the model also tracks various metrics like CPEFH, mission completion rate, etc.

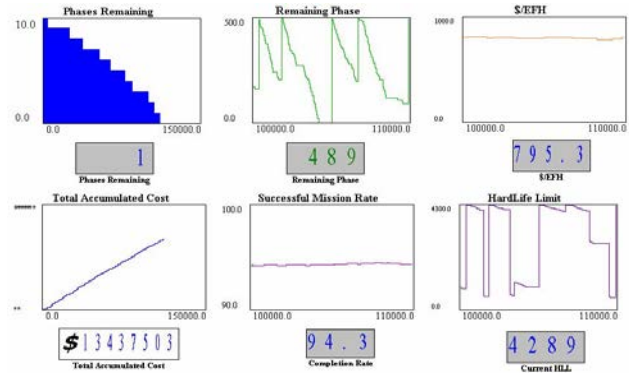


Figure 4: Cost Projection Simulation Outputs

This simulation framework provides the capability to determine LCC based on reliability characteristic curves. More importantly, the simulator provides a framework to analyze how supply chain parameters (on-hand inventories, due-ins, local inventories, engines available for swapping) and maintenance policy options (whether to repair, swap with another module, replace with a new module, or defer maintenance) affect LCC.

This model functionality can also support the generation of the Cost per Mean Time Between Removals chart that is a common reference for thinking about the characteristic tradeoff between cost and operational availability. The ability to generate this curve will allow further study of its characteristic form for different type, model, and series (TMS) engines operated under various maintenance policies.

The prototype CPS developed as part of this research demonstrated the feasibility of the concept of using a discrete-event engine to develop estimates of LCC. The model design accounts for various factors, including environmental stresses (e.g., fine sand environments) and changes in op tempo. These factors are used to adapt individual component Weibull parameters over time as the engines are flown in the simulator.

6.2 Data Mining Results

Data mining refers to a loose collection of analytical driven techniques and technologies that support the overall knowledge discovery process (Benjamin 2002). KBSI's Personal Data Prospector (PDP) was used to perform data mining on the simulation output. The PDP includes an extensive set of data access, manipulation, mining, and visualization functions supported by fundamental data manipulation elements, data processing elements, numerical algorithmic components, and data visualization components. These elements can be strung together graphically to define and execute an analytical process. In this way, the data mining process logic is captured explicitly as a process diagram or network (called *Knowledge Analysis Net-*

works). This mechanism makes the data mining process transparent and repeatable.

The focus of this data mining effort was to discover and understand LCC drivers. Thus, the approach used extracts engine life-cycle events in terms of projected maintenance issues, actions performed, the subsequent events that resulted, and their impact in terms of cost. In particular, this approach uses regression, clustering, and classification to determine cost drivers. The following sections summarize the results of those efforts.

6.2.1 Linear Regression

Regression analysis models the relationships between TOW or more explanatory variables (or independent variables) and a response variable (dependent variable). In linear regression, the relationship is modeled as a linear equation. The most popular approaches to linear regression attempt to fit the best line to the data by minimizing the mean square errors (square of deviation from the data point and the line).

Formally, the model for linear regression can be modeled as:

$$\text{Data} = \text{Linear Model} + \text{Residual},$$

$$Y = \text{Constant} + \alpha_1 X_1 + \alpha_2 X_2 + \dots + \alpha_n X_n + \varepsilon,$$

where:

Y is the dependent variable,
 X_1, X_2, \dots, X_n are explanatory variables, and
 ε is the residual.

One of the first activities undertaken to study the simulation data was to develop a linear regression model to understand the variables influencing LCC. LCC was modeled as a dependent, or response, variable. The number of repairs, swaps with inventory, swaps with modules that were still part of other engines, replacements with new modules, and the number of no actions (i.e., deferred maintenance decisions) were modeled as independent or explanatory variables.

Regression analysis resulted in the following model:

$$\text{LCC} = (450,556 + 61,451 * \text{Number Repairs} + 158,430 * \text{Number of New Replacements} + 18,766 * \text{Number of Swaps Between Engines} + 55,471 * \text{Number of Swaps with Inventory} + 107,042 * \text{Number of No Actions})$$

The regression model had a very high F value (346) indicating at least one linearly dependent variable. The model had a reasonable R-square (.77) and Adjusted R-square (.77) indicating that the model was able to explain LCC variability to a reasonable extent.

The following intuitive notions were obtained through an analysis of the regression model:

1. The most significant factor affecting LCC was the number of new replacements. This variable had a coefficient of 158,430.
2. Another significant factor affecting LCC was the number of repairs. This variable had a coefficient of 61,451.

Several counter-intuitions were also identified. Validation of the model and data results is needed to make sure that these counter-intuitive notions are not a result of approximations in the model. If validated, these observations can yield good insights into cost drivers:

3. The coefficient for swapping modules between engines (18,766) was less than the coefficient of swapping modules with assets in inventory (55,471). We expected the cost of swapping between engines to be higher because of the overhead involved. One possible explanation for this could be that swapping between engines resulted in better-aligned engines.
4. The coefficient for 'not taking any action' was unusually high (107,042). This could be due to the modules failing later on, thereby affecting the mission.

6.2.2 Classification Methods

Classification methods refer to a collection of statistical or heuristic methods that enable the separation of a set of entities (characterized by a representative attributes) into groups (classes) and allocating new entities into these pre-defined groups (classes). The final result is a set of rules that help in assigning new observations to one of the classes. Another way of looking at classification is as a partitioning of the independent variable space (or attribute space) into regions, where each region is associated with a particular class assignment. The boundaries between these regions are called decision surfaces.

Commonly used classification methods include logistic regression models, multiple discriminant analysis (MDA), decisions trees for classification (CARD and CHAID), artificial neural networks (ANN) for classification, and Bayesian networks.

Classification techniques were used to analyze the parameters that influence LCC. All data points were classified as either low- or high-cost engines based on their LCC. Classification was done using the CART algorithm and cost classes were used to develop the classification tree structure. The variables — number of repairs, new module replacements, swaps between engines, swaps with inventory, number of repairs within 50 hours of phase (planned

intervals of running-hours that define a maintenance point), number of new replacements within 100 hours of phase, and number of no actions — were chosen as attributes to be included in the model to develop the classification tree (Figure 5).




Figure 5: Classification Tree

For validation purposes, the classification tree was used to classify the original data points into low- or high-cost classes. The classifier performed well in accurately predicting the cost class based on such attributes as number of new replacements, number of repairs, etc. This validated that the classifier developed was reasonable in identifying LCC drivers.

Several insights were gained through the classification exercise, including the following.

5. Repairing parts was the most statistically significant factor influencing LCC. The classifier picked number of repairs as the first attribute to distinguish between low- and high-cost data points. This could be caused by incurring significant labor and material cost through repair without significantly improving asset life expectancy. This insight may be somewhat counter-intuitive, but very significant.
6. The second most significant variable influencing costs involved replacement with new modules. When the number of repairs was greater than 9.5 and number of replacements was greater than 6.5, the LCC of an engine would automatically be high. If the number of new replacements was greater than 11.5, LCC was high irrespective of anything else.
7. While it might be intuitive to expect that LCC is high for a high number of replacements with new modules, it is not obvious that performing too many repairs can also increase LCC. This finding is very significant and reinforces two observations:
 - (a) Over-performing repair activities can be counter-productive in some situations.
 - (b) For the specific situation simulated, swapping may be a better option than repairing.

6.2.3 Clustering Techniques

Clustering is an exploration of a data set to determine the structure of natural groupings. Clustering techniques are distinct from classification techniques since they pertain to a known number of groups. While the objective in classification is to assign new observations to one of these groups, cluster analysis is more primitive in that no assumption is made about the number of underlying groups or any other structure. Grouping is done after defining an appropriate similarity or distance measure. The objective of optimization-based clustering is to determine a partitioning of the data set that minimizes intra-cluster variability and/or maximizes inter-cluster variability. To this end, a clustering criterion (i.e., objective in an optimization model) is defined, and the given samples are classified to optimize this criterion. Minimization of the intra-cluster variability and maximizing the inter-cluster variability yields the cluster centers along with the assignments of the samples to the different clusters. Popular methods include hierarchical clustering methods (agglomerative hierarchical method and divisive hierarchical methods) and neural net-based technologies such as Kohonen self-organizing feature maps.

The simulation results generated by the CPS were analyzed using the clustering data mining technique. The focus of clustering efforts was to segment low-, medium- and high-LCC engines and then study them to understand the variables or factors that influenced their costs.

The data elements used for clustering include the LCC of the engine, the number of swaps with inventory, swaps with engines, number of new replacements, number of repairs, and number of no actions. Since we wanted to analyze whether the timing of performing a maintenance action influences LCC, we included two additional variables: (1) New Replacements within 100 hours of Phase, and (2) Repair within 50 hours of Phase.

Clustering was done on a single variable (LCC) using the K-means algorithm. All the other variables were used for profiling. Profiling involves calculating the mean value of different variables in each cluster to study the similarities and differences between clusters.

In this study, the data elements were grouped into three clusters to represent low-, medium-, and high-cost data points. After the clusters were formed, a profile was developed for each cluster by calculating the mean values of the explanatory variables for each cluster. The centroids of the three clusters are summarized in Table 1.

Table 1. Profile Results for Each Cluster

Attribute	Cluster 1	Cluster 2	Cluster 3
CumulativeCost_Mean	\$3,390,322	\$3,822,992	\$4,269,430
RecordCount	139	249	112
ReplaceNew_Mean	7.97	9.84	10.69
Repair_Mean	7.24	8.93	11.69
RepairWithin50HPhase_Mean	0.68	0.71	0.95
ReplaceNewWithin100Phase_Mean	1.32	1.45	1.37
SwapBEngs_Mean	9.73	8.33	7.32
SwapWithin_Mean	9.15	9.65	11.69
NoRepair_Mean	6.01	5.19	5.02

The following insights were gained through the clustering exercise:

8. Low-, medium-, and high-cost clusters (Clusters 1, 2 and 3 in Table 1) differ significantly in their Replace New mean (8.0, 9.8, and 10.7; respectively) and Repair mean (7.2, 8.9, and 11.7; respectively). This result reinforces the insights gained through regression analysis.
9. Replacing new engine modules within 100 hours of phase did not seem to influence the cost. All three clusters (low-, medium- and high-cost clusters) were very similar in terms of the number of replacements of new engines within 100 hours of Phase (1.32, 1.45, and 1.37; respectively).
10. Repairing a module within 50 hours of phase seems to have a slight effect on cost differences. The mean values of the three clusters were 0.68, 0.71 and 0.95. Further study would be needed to verify whether this difference is statistically significant or not.

7 SUMMARY AND CONCLUSIONS

A LCC estimation methodology supported by discrete event simulation and data mining provides a robust approach that overcomes practical issues that cannot be addressed using traditional methods. This innovation provides a key capability supporting RCM-based decision-making.

Demonstration of concept viability was accomplished by first showing that discrete event simulation can be used to project LCC. A distinct advantage of a discrete event simulation-based approach is that LCCs are determined by considering not only failure rates and Weibull characteristic curves, but a number of important factors that influence LCC — like inventory levels, swapping policies, supply posture, the environment of use, increases in part acquisition costs, etc. Coupled with this was a demonstration of the ability to use data mining to identify and characterize cost drivers by extracting them from simulation data output.

These findings represent significant innovations that, when fully developed, can help both field and depot engine maintenance units successfully apply RCM principles and maximize the return in engine flight hours for each dollar

spent toward maintenance. Significant developments from this research include the following:

- Clear identification of the theoretical and practical issues involved in LCC projection.
- Unique approach coupling discrete event simulation and data mining technologies to either directly calculate or develop and maintain parameters to determine the LCC implications of maintenance decisions made in a given context. This approach circumvents the risks and shortcomings of parametric model-based approaches, minimizes dependence on maintenance history and accounting data, directly leverages models reflecting component reliability characteristics, and accounts for complex interactions among constraints in the operational environment to provide statistically rigorous and accurate LCC projections.
- Simulation-based capability to project the cost dynamics and overall LCC implications of maintenance decisions, thus providing a significant advancement in LCC estimation capabilities.

ACKNOWLEDGMENTS

This research was funded by the Air Force Research Laboratory, Propulsion Directorate (AFRL/PRTS), at Wright-Patterson Air Force Base, Ohio under the technical direction of Dr. Brian Beachkofski. Assisting him in this role were Mr. Gary Smith of AFRL/PRTS and Ms. Cheryl Holtz of Universal Technology Corporation (UTC). The authors would also like to acknowledge the contributions of Mr. Scott Lasiter, the Hawk RCM program lead at Oklahoma City Air Logistics Center (OC-ALC), whose technical insights, vision, and support were pivotal to the success of this effort. Mr. Tom Felux, the Air Force TF39 RCM program lead, also provided invaluable technical input reflecting his work in estimating engine remaining life characteristics.

REFERENCES

- Beachkofski, B. 2006. "Pervasive Panel: Safety and Readiness," *P-SAR Conference proceeding*. Air Force Research Laboratory.
- Benjamin, P. C., M. Erraguntla, R. J. Mayer. 1995. Using Simulation for Robust System Design. *Simulation* 65(2), pp. 116-128.
- Benjamin, P., E. Erraguntla, S. Ramachandran, and D. Delen. 2002. Towards a Knowledge Discovery Framework. In *Proceedings of the 2002 International Conference on Information and Knowledge Engineering (IKE 2002)*. Las Vegas, June 2002. CSREA Press. Pages 16-22.

- Cifone, A. 2006. Propulsion Safety & Affordable Readiness (P-SAR): A View from the Navy. *P-SAR Conference Proceedings*, Air Force Research Laboratory.
- CNAF. 2006. Aviation safety standdown orders. Commander Naval Air Forces. San Diego, CA. Available at <http://www.safetycenter.navy.mil/aviation/downloads/040104ZMAR06.txt>.
- Erraguntla, M., P. C. Benjamin, and R. J. Mayer. 1994. An Architecture of a Knowledge-Based Simulation Engine. In *Proceedings of the 2004 Winter Simulation Conference*, pp. 673-680. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.
- Goldberg, D. E. 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, MA.
- Nowlan, F. S., and H. F. Heap. 1978. *Reliability-Centered Maintenance*. United Airlines and Dolby Press. Sponsored and published by the Office of Assistant Secretary of Defense.
- Painter, M. K., M. Erraguntla, and G. Hogg. Reliability Centered Maintenance Costing (RCMC) Final Report, Air Force SBIR Contract No. FA8650-05-M-2568.
- SAE Standard JA 1011. 1999. Evaluation Criteria for Reliability-Centered Maintenance (RCM) Processes.

AUTHOR BIOGRAPHIES

MICHAEL K. PAINTER is a Program Manager and leading researcher at Knowledge Based Systems, Inc. He received a B.S. in Mechanical Engineering with a minor in Japanese from Utah State University and a Masters in Business Administration from Texas A&M University. He manages advanced solutions development projects, providing technical direction and leadership to teams of research scientists, operations research analysts, simulation analysts, knowledge engineers, data miners, process modelers, systems architects, computer scientists, and transformation consultants.

DR. MADHAV ERRAGUNTILA received his Ph.D. in Industrial Engineering from Texas A&M University in 1996 and Master's in Industrial Engineering from the National Institute for Training in Industrial Engineering in 1989. He is a research scientist at KBSI and has conducted extensive research on and developed applications in knowledge representation, ontology, text mining, data mining, optimization, agent based systems, simulation, evolutionary computing, analytics, planning, activity based costing, knowledge-based systems, neural networks, and fuzzy logic.

GARY L. HOGG, JR. is a program manager and researcher at Knowledge Based Systems, Inc. He received

his Masters in Industrial Engineering from Texas A&M University. He is a researcher in the analysis of advanced manufacturing processes and systems and has an extensive background in systems modeling using simulation and analytic techniques, optimization methodologies, and queueing theory. Mr. Hogg currently supports projects in improving maintenance, repair, and overhaul (MRO) processes for military rotary-wing aircraft, engines, transmissions and components at the Corpus Christi Army Depot.

BRIAN BEACHKOFSKI is a distinguished graduate of the US Air Force Academy class of 1998 where he earned his BS in Engineering Mechanics. He followed that by earning his MS in aerospace engineering at Georgia Tech, graduating in 1999. His Ph.D. in Mechanical Engineering from Wright State University was completed in 2004, focusing on probabilistic analysis methods to reduced confidence intervals. He is currently with the Air Force Research Laboratory's Turbine Engine Division as the engine health management lead.