

A Decision Support Model for Determining the Applicability of Prognostic Health Management (PHM) Approaches to Electronic Systems

Peter Sandborn, CALCE, Dept. of Mechanical Eng., University of Maryland

Key Words: prognostic health management (PHM), decision support, cost modeling, imperfect monitoring

SUMMARY & CONCLUSIONS

This paper presents a model that enables the determination of when scheduled maintenance makes sense, and how to optimally interpret Prognostic Health Management (PHM) results for electronic systems. In this context, optimal interpretation of PHM results means translating PHM information into maintenance policies that minimize life cycle costs. The electronics PHM problem is characterized by imperfect and partial monitoring and a significant random/overstress failure component must be considered in the decision process. Specifically the model enables determining on an application-specific basis when the reliability of electronics has become predictable enough to warrant the application of PHM-based scheduled maintenance concepts. Given that the forecasting ability of PHM (whether health monitoring or life consumption monitoring based) is fraught with uncertainties in the sensor data collected, the models applied, the material parameters assumed in the models, etc., the model in this paper addresses how PHM results can be interpreted so as to provide value to the system. The result of the model is the determination of optimal safety margins on life consumption monitoring predictions and prognostic distances for health monitoring.

The model provides the type of information needed to construct a business case showing the application-specific usefulness of health monitoring and/or life consumption monitoring for electronic systems.

1. INTRODUCTION

Prognostics is the estimation of remaining life in terms that are useful to the maintenance decision process. All PHM approaches are essentially the extrapolation of trends based on recent observations to estimate remaining life, [1]. Unfortunately, this calculation alone does not provide sufficient information to form a decision or to determine corrective action. Without comprehending the corresponding measures of the uncertainty associated with the calculation, remaining life projections have little practical value, [1]. It is the comprehension of the corresponding uncertainties (decision making under uncertainty) that is at the heart of being able to develop a business case that addresses prognostic requirements.

Electronic systems have not traditionally been subject to Prognostic Health Management (PHM) because their time to

wear-out was assumed to be much longer than the system life cycle or technology refresh period (non-life limited). Most approaches to PHM are focused on monitoring failure precursor indications (i.e., health monitoring), which does not require system failures to be deterministic in nature, but does require that the failure precursor have a deterministic link to the actual system failure. While there is considerable existing work on failure precursors for mechanical systems, only a few attempts have been made to apply health monitoring to electronics, [2-3]. Alternatively, Life Consumption Monitoring (LCM), which is another approach to PHM, depends on the deterministic nature of system failures. In LCM, a history of environmental stresses (e.g., thermal, vibration) is used in conjunction with physics of failure models to compute damage accumulated and thereby forecast life remaining, [4]. Electronic systems have not traditionally been subject to LCM because the distribution of failures prior to wear-out was considered to be random (non-deterministic). With the transition from military-specification parts to commercial-off-the-shelf (COTS) parts, many of which are now targeted for lifetimes in the 5 to 7 year range, wear-out of electronics parts is becoming a relevant concern for long field life systems, [5]. In addition, physics of failure approaches to modeling electronic system reliability have shown that time-to-failure prior to wear-out for electronic parts can be predicted within quantifiable bounds of uncertainty, [6].

Modeling to determine the optimum schedule for performing maintenance for systems is not a new concept. Examples of traditional applications of maintenance modeling include production equipment [7] and the hardware portions of engines and other propulsion systems [8]. However, maintenance modeling has not been widely applied to electronic systems where presumed random electronics failure is usually modeled as an unscheduled maintenance activity, and wear-out is beyond the end of the system's life.

Although many applicable models for single and multi-unit maintenance planning have appeared [9,10], the majority of the models assume that monitoring information is perfect (without uncertainty) and complete (all units are monitored the same), i.e., maintenance planning can be performed with perfect knowledge as to the state of each unit. For many types of systems, and especially electronic systems these are not good assumptions and maintenance planning, if possible at all, becomes an exercise in decision making under uncertainty with sparse data. The perfect monitoring assumption is especially problematic when the PHM approach is Life

Consumption Monitoring (LCM) because LCM does not depend on detecting precursors to failure. When managing electronic systems, system-level and major component failures are caused by a mixture of failure mechanisms. These failure mechanisms result from defects, wearout, overstress conditions, and random system interactions. These types of failures are a mixture of predictable and partially predictable events. Thus, for electronics, LCM PHM processes do not deliver any measures that correspond exactly to the state of the system. Previous work that treats imperfect monitoring includes, [11] and [12]. Perfect, but partial monitoring has been treated in [13].

This paper presents a model which determines when scheduled maintenance makes good business sense. The model shows how to optimally interpret damage accumulation or failure precursor monitoring data. This will apply to failure events that appear to be random or appear to be clearly caused by defects, wearout, or overstress conditions. This optimal interpretation of the data means that we can optimize system availability and minimize life cycle cost. Specifically the model is targeted at addressing the following questions:

- How do we determine on an application-specific basis when the reliability of electronics has become predictable enough to warrant the application of PHM-based scheduled maintenance concepts? Note, we do not mean to imply that predictability in isolation is the criteria for PHM vs. non-PHM solutions, e.g., if the system reliability is predictable and very reliable, it would not make sense to implement a PHM solution.
- Given that the forecasting ability of PHM (health monitoring or life consumption monitoring based) is fraught with uncertainties in the sensor data collected, the data reduction methods, the models applied, the material parameters assumed in the models, etc., how can PHM results be interpreted so as to provide value? This boils down to determining optimal safety margins on life consumption monitoring predictions and prognostic distances¹ for health monitoring.
- How can a business case be constructed to show the usefulness of health monitoring and/or life consumption monitoring for electronic systems?

This paper describes the formulation of a single LRU (Line Replaceable Unit) stochastic optimization model.

2. MODEL FORMULATION

The following model accommodates variable time-to-failure and LCM forecast distributions. The model considers only one LRU (Line Replaceable Unit) within a larger system. The model treats all inputs as probability distributions, i.e., a stochastic analysis is used (Monte Carlo). Various maintenance interval and PHM approaches are distinguished by how sampled time-to-failure values are used to model PHM forecasting distributions. The metrics computed are: life cycle

¹ The duration (e.g., measured in time or cycles) between the actual failure and the point where the prognostic structure fails or indicates failure.

cost, failures avoided, and operational availability. Nothing about the modeling is electronic system specific.

To assess PHM, relevant failure mechanisms must be segregated into two types:

- Failure mechanisms that are random from the view point of the PHM methodology. These are failure mechanisms that the PHM methodology is not collecting any information about (non-detection events). These failure mechanisms may be predictable but are outside the scope of the PHM methods applied.
- Failure mechanisms that are predictable to some degree from the view point of the PHM methodology, i.e., for which a probability distribution can be assigned.

Several cases are considered in the model that follows: 1) a fixed scheduled maintenance interval that is kept constant for all instances of the LRU throughout the system life cycle; where an “instance” is one particular fielded LRU; 2) a variable maintenance schedule for the LRU that is based on inputs from a Life Consumption Monitoring (LCM) methodology; and 3) a variable maintenance scheduled for the LRU that is based on a Health Monitoring (HM) methodology. Note, the model formulation is presented based on “time” to failure measured in operational hours, however, the relevant quantity could be a non-time measure such as thermal cycles.

Example results generated using all the approaches discussed in this section are presented in Section 3.

2.1 Fixed Scheduled Maintenance Interval

This case is well understood, but included herein because it serves to define the general approach that will be used for assessing health monitoring and life consumption monitoring. In this case a fixed scheduled maintenance interval is selected that is kept constant for all instances of the LRU throughout the system life cycle. In this case the LRU is replaced on a fixed interval (measured in operational hours), i.e., time-based prognostics. Consider the simply represented time to failure distribution shown in Fig. 1.

The following algorithm is used to accumulate life cycle costs (C) based on time stepping through the lifetimes of a statistically relevant set of LRUs where T is time:

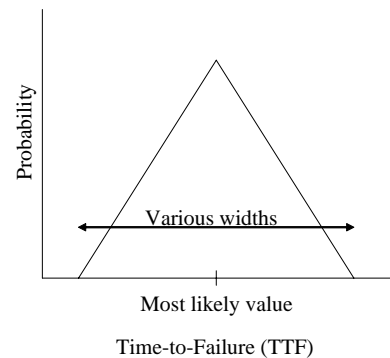


Fig. 1. Symmetric triangular time to failure distribution.

Note, the model is not constrained in any way to working with either symmetric or triangular distributions. Other distributions can be used.

1. Defined Time-to-Failure (TTF) distributions (subscript R = random, subscript P = predictable)
2. Sample the TTF distributions to get TTF_R and TTF_P
3. Compare $TTF_S = \min(TTF_R, TTF_P)$ to a defined fixed maintenance interval (MI)
4. If $TTF_S \leq MI$ then $T = T + TTF_S$ and $C = C + C_{us}$ (subscript us = unscheduled)
5. If $TTF_S > MI$ then $T = T + MI$ and $C = C + C_s$ (subscript s = scheduled) – in this case a failure was avoided
6. Repeat steps 2-5 until $T >$ operation and support life of the system
7. Repeat steps 1-6 a statistically relevant number of times in order to build histograms of life cycle costs, availability, and failures avoided.

The random times-to-failures, TTF_R , are characterized by an average random failure rate per operational period (e.g., 1 year) expressed as a fraction of the total units fielded. To model the random failure rate, a uniform distribution with a height equal to the average random failure rate per year and a width equal to the inverse of the average random failure rate is created and sampled to get TTF_R .

2.2 Life Consumption Monitoring (LCM)

Life Consumption Monitoring is defined in this paper as the process by which a history of environmental stresses (e.g., thermal, vibration) is used in conjunction with physics of failure models to compute damage accumulated and thereby forecast life remaining.

The LCM methodology forecasts a unique time-to-failure distribution for each instance of the LRU based on its unique environmental stress history. For this example, the LCM forecast was modeled as a symmetric triangular distribution with a most likely value set to the time-to-failure of the LRU instance and a fixed width measured in operational hours, (Fig. 2). The shape and width of LCM distribution depends on the uncertainties associated with the sensing technologies and uncertainties in the prediction of the damage accumulated (data and model uncertainty). The variable to be optimized in this case is the safety margin assumed on the LCM forecasted time-to-failure, i.e., the length of time (e.g., in operation hours) before the LCM forecasted failure the unit should be replaced. The model proceeds in the following way: for each

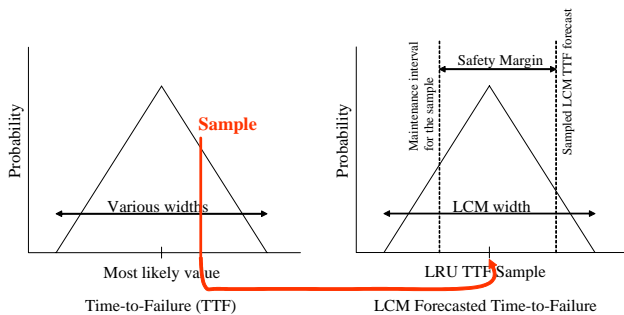


Fig. 2. Life consumption monitoring modeling approach. Symmetric triangular distributions are shown for simplicity.

time-to-failure distribution sample, an LCM distribution is created that is centered on the time-to-failure. The LCM distribution is then sampled and if the LCM sample minus the safety margin is less than the actual time-to-failure of the LRU instance then LCM was successful (failure avoided). If the LCM distribution sample minus the safety margin is greater than the actual time-to-failure of the LRU instance then LCM was unsuccessful. If successful, a scheduled maintenance activity is performed and the timeline is incremented by the LCM sampled time-to-failure minus the safety factor. If unsuccessful, an unscheduled maintenance activity is performed and the timeline is incremented by the actual time-to-failure of the LRU instance. As with the fixed scheduled maintenance interval approach, the LCM model is implemented as a stochastic simulation. In this simulation, a statistically relevant number of LRUs are considered in order to construct Histograms. The Histograms allow us to choose safety margins that minimize mean life cycle cost or alternatively minimize the cumulative life cycle cost of all units sampled. Also similar to the fixed interval case, a random failure component is superimposed.

2.3 Health Monitoring (HM)

Health monitoring is defined in this paper as monitoring for failure precursors. For health monitoring the methodology presented in this paper can be used to determine prognostic distance. Where the prognostic distance is a measure of how long before system failure the prognostic structures or prognostic cell is expected to fail or indicate failure (in operational hours for example). The health monitoring methodology forecasts a unique time-to-failure distribution for each instance of the LRU based on its time-to-failure. For this example, the health monitoring forecast was modeled as a symmetric triangular distribution with a most likely value set to the time-to-failure of the LRU instance minus the prognostic distance, Fig. 3. The health monitoring distribution has a fixed width measured in the relevant environmental stress units (e.g., operational hours in our example) representing the probability of the prognostic structure indicating the precursor to a failure. As a simple example, if the prognostic structure was a fuse that is designed to fail at some prognostic distance earlier than the system it protects, then for this example the distribution on the right side of Fig. 3 represents the distribution of fuse failures. The variable to be

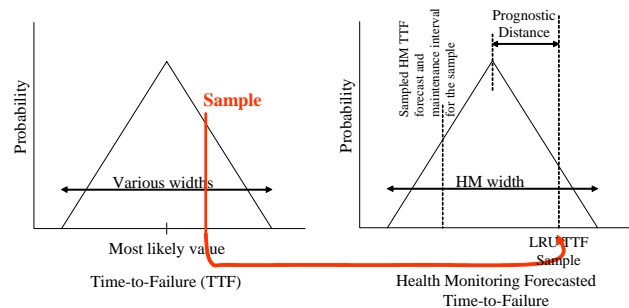


Fig. 3. Health monitoring modeling approach. Symmetric triangular distributions are shown for simplicity.

optimized in the HM case is the prognostic distance assumed on the health monitoring forecasted time-to-failure. The model proceeds in the following way: for each time-to-failure distribution sample, a health monitoring distribution is created that is centered on the time-to-failure minus the prognostic distance. The health monitoring distribution is then sampled and if the health monitoring sample is less than the actual time-to-failure of the LRU instance then health monitoring was successful. If the health monitoring distribution sample is greater than the actual time-to-failure of the LRU instance then health monitoring was unsuccessful. If successful, a scheduled maintenance activity is performed and the timeline is incremented by the health monitoring sampled time-to-failure. If unsuccessful, an unscheduled maintenance activity is performed and the timeline is incremented by the actual time-to-failure of the LRU instance.

For both the LCM and HM approaches, the performance of the PHM methodology is modeled as a probability distribution taking into account uncertainties embedded in the methodology, sensors, models, etc. Note that in both cases the PHM probability distributions are coupled to the actual failure of an instance of the LRU.

3. MODEL RESULTS

Variables in the current model are shown in Table I.

Table I. Data assumptions for cases presented in this paper.

Variable in the model	Value used for example analysis	
Production cost (per unit)	\$10,000	
Time to failure	5000 operational hours = the most likely value (symmetric triangular distribution with variable distribution width)	
Operational hours per year	2500	
Sustainment life	25 years	
	Unscheduled	Scheduled
Value of each hour out of service	\$10,000	\$500
Time to repair	6 hours	4 hours
Time to replace	1 hour	0.7 hours
Cost of repair (materials cost)	\$500	\$350
Fraction of repairs requiring replacement of the LRU (as opposed to repair of the LRU)	1.0	0.7

All of the variable inputs to the model can be treated as probability distributions or as fixed values, however, for example purposes, only the time to failure and performance of the PHM methodologies have been characterized by probability distributions. Note, all the life cycle cost results provided in the remainder of this paper are the mean life cycle cost from a probability distribution of life cycle costs generated by the model.

Figure 4 shows the fixed scheduled maintenance interval results. 10,000 LRUs were simulated in a Monte Carlo analysis and the mean life cycle costs are plotted in Fig. 4. The general characteristics in Fig. 4 are intuitive: For short scheduled maintenance intervals, virtually no expensive unscheduled maintenance occurs, but the life cycle cost per unit is high because large amounts of remaining life in the LRU are thrown away. For long scheduled maintenance intervals virtually every instance of the LRU fails prior to the scheduled maintenance activity and the life cycle cost per unit becomes equivalent to a simple unscheduled maintenance model. For some scheduled maintenance interval between the extremes, the life cycle cost per unit is minimized. If the time-to-failure distribution had a width of zero, then the optimum fixed scheduled maintenance interval would be exactly equal to the time-to-failure. As the time-to-failure distribution becomes wider (i.e., the time-to-failure is less well defined), a practical fixed scheduled maintenance interval becomes more difficult to find and the best solution approaches an unscheduled maintenance model.

Figure 5 shows example results from life consumption monitoring. Several general trends are apparent. First, the width of the time-to-failure distribution has little effect on the LCM-based results. It is important to note that in this model,

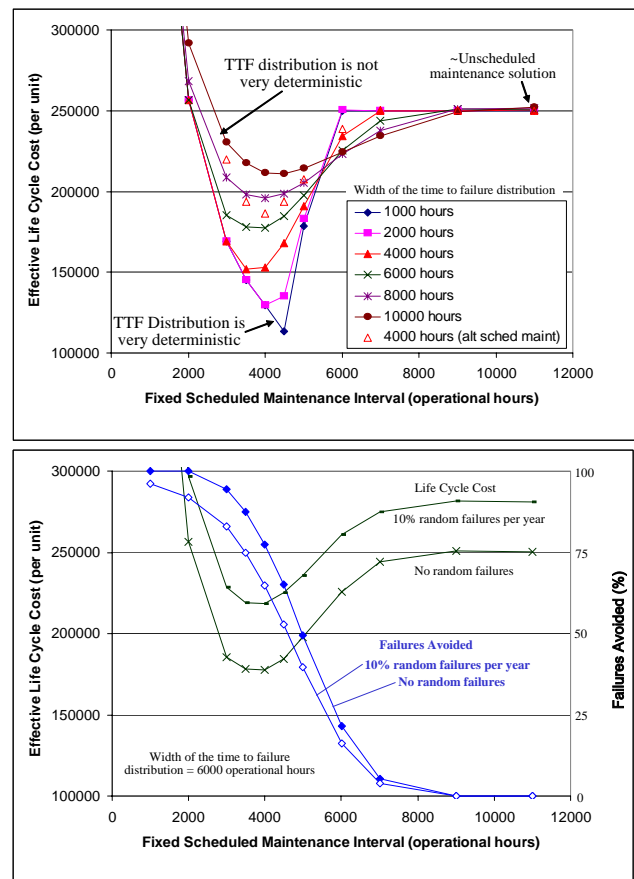


Fig. 4. Variation of the effective life cycle cost per unit with the fixed scheduled maintenance interval (10,000 Monte Carlo samples). Top: no random failures assumed; Bottom: 10% random failures per year included, variation in failures avoided also shown.

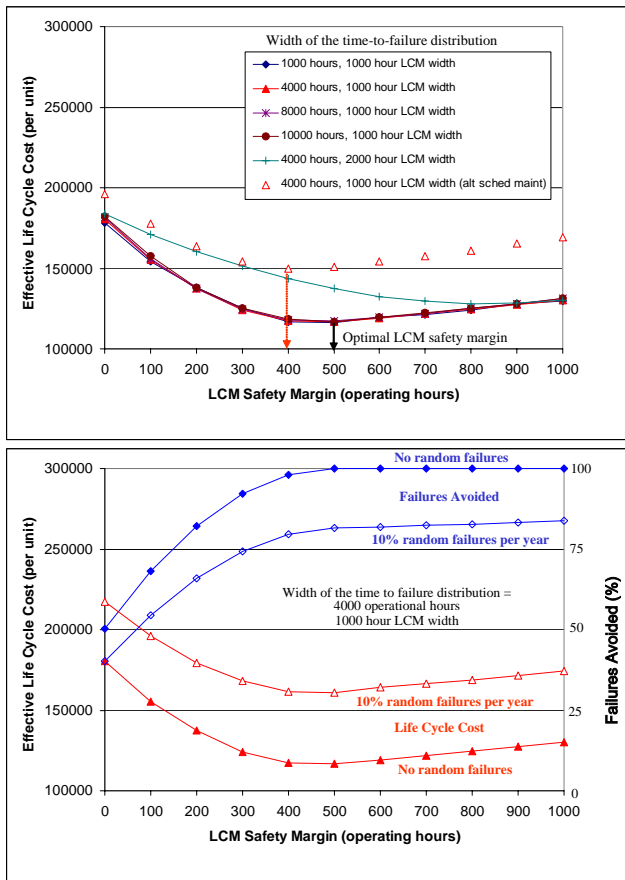


Fig. 5. Variation of the effective life cycle cost per unit with the safety margin for a LCM-based maintenance-planning scheme (10,000 Monte Carlo samples). Top: no random failures assumed; Bottom: 10% random failures per year included, variation in failures avoided also shown.

the characteristics of the actual time-to-failure distribution are uncoupled from the characteristics of the LCM forecasted time-to-failure distribution, i.e., they are assumed to have different widths. This is reasonable because the actual time-to-failure distribution is a distribution of the time-to-failures of a large population of LRUs, and the LCM forecasted time-to-failure distribution is the prediction of the time-to-failure of a specific LRU given uncertainties in its environmental stress history, damage accumulated prior to the start of life consumption monitoring, and variations in its materials and manufacturing. The second trend is that for all but the narrowest distribution on the LRU time-to-failure, the LCM-based methodology yields lower life cycle costs than either unscheduled maintenance or scheduled maintenance with a fixed interval. However, for a very narrow time-to-failure distribution, a fixed scheduled maintenance interval may yield a lower life cycle cost solution than a LCM-based methodology. We have also found that the effective life cycle cost is very sensitive to the width of the LCM forecasted time-to-failure distribution for the specific LRU (one case with a 2000 hour width is shown in Fig. 5).

Figure 6 shows example results from a health monitoring solution for the example described in Table I. Figure 6




Fig. 6. Variation of the effective life cycle cost per unit with the prognostic distance for a health monitoring based maintenance planning scheme (10,000 Monte Carlo samples). Top: no random failures assumed; Bottom: 10% random failures per year included, variation in failures avoided also shown.

indicates that the health monitoring based methodology yields lower life cycle costs than either unscheduled maintenance or scheduled maintenance with a fixed interval. However, life consumption monitoring results in lower life cycle costs (at least for the numbers assumed in this example). Figure 6 also shows that the effective life cycle cost is very sensitive to the width of the health monitoring forecasted time-to-failure distribution for the specific LRU.

Figures 4-6 include plots of the failures avoided for several of the cases considered with and without random failures included in the analysis. In all cases the failure avoided when random failures are included is lower than when random failures are not included and the safety margin or prognostic distance is small (or fixed scheduled maintenance interval is large). As the safety margin or prognostic distance increase (or fixed scheduled maintenance interval is small) the failures avoided limits to 100% in all cases (with and without random failures included). However, for the example data used in this paper, safety margins or prognostic distances must be increased substantially beyond the range plotted in Figs. 4-6 for the cases with random failures to tend to 100%.

4. DISCUSSION

Previous PHM work on electronic systems has demonstrated life consumption monitoring for electronic systems, [4]. However, the previous work has not addressed if (or exactly how) a specific application's life cycle cost can actually be reduced and/or operational availability increased by using PHM. Such an analysis becomes non-trivial when one considers the accuracy that life consumption monitoring results are likely to have (imperfect and partial monitoring conditions). Schemes for interpreting and applying PHM results to maintenance decisions will have to balance the risk of unscheduled failure with the substantial uncertainties present in the life consumption monitoring results.

The 1990's perfected obtaining and storing large amounts of information, and as a result, the world is wading in a lot more information that it knows how to use. The trick now is to figure out how to make decisions based on the information – particularly when the information is incomplete and uncertain.

The single LRU model presented in this paper is being extended to treat multiple LRUs. In addition, several other key questions ultimately need to be addressed including:

- What is the right shape and size of distributions associated with Life Consumption Monitoring and Health Monitoring?
- Realistic PHM solutions for electronic systems will probably be mixtures of LCM, HM and scheduled maintenance
- Treatment of redundancy
- Second order uncertainty (uncertainty about uncertainty) may be a real issue in the treatment of this problem.

REFERENCES

1. S. Engel, B. Gilmartin, K. Bongort, A. Hess, "Prognostics, the Real Issues Involved with Predicting Life Remaining," *Proceedings of the IEEE Aerospace Conference*, 2000, pp. 457-469.
2. N. Kelkar, A. Dasgupta, M. Pecht, I. Knowles, M. Hawley D. Jennings, "'Smart' Electronic Systems for Condition-Based Health Management," *Qual. and Reliability Eng. Int.*, Vol. 13, 1997, pp. 3-7.
3. S. Mishra, M. Pecht, D. Goodman, "In-situ Sensors for Product Reliability Monitoring," *Proceedings of SPIE*, Vol. 4755, 2002, pp. 10-19.
4. A. Ramakrishnan, M. Pecht, "A Life Consumption Monitoring Methodology for Electronic Systems," *IEEE Transactions on Components and Packaging Technologies*, (September) 2003, pp. 625-634.
5. D. Humphrey, D. Lorenson, W. Shawlee, P. Sandborn, "Aging Aircraft Usable Life and Wear-out Issues," *Proc. of the World Aviation Congress (SAE Tech. Paper: 2002-1-3013)*, Phoenix, AZ, (November) 2002.
6. *Product Reliability, Maintainability, and Supportability*

Handbook, edited by M. Pecht, CRC Press, New York, NY, 1995.

7. M. Ben-Daya, S. O. Duffuaa, A. Raouf, *Maintenance, Modeling and Optimization*, Kluwer Academic Publishers, Boston, MA, 2000.
8. G. J. Kacprzynski, M. Gumina, M. J. Roemer, D. E. Caguiat, T. R. Galie, J. J. McGroarty, "A Prognostic Modeling Approach for Predicting Recurring Maintenance for Shipboard Propulsion Systems," *Proc. of ASME Turbo Expo*, (June) 2001.
9. C. Valdez-Flores, R. Feldman, "A Survey of Preventative Maintenance Models for Stochastically Determining Single-Unit Systems," *Naval Research Logistics*, Vol. 36, 1989, pp. 419-446.
10. D. Cho, M. Parlar, "A Survey of Preventative Maintenance Models for Multi-Unit Systems," *European J. Operational Research*, Vol. 51, 1991, pp. 1-23.
11. W. Wang, "A Model to Determine the Optimal Critical Level and the Monitoring Intervals in Condition-Based Maintenance," *Int. J. Production Research*, Vol. 38, No. 6, 2000, pp. 1425-1436.
12. A. Barros, C. Berenguer, A. Grall, "Optimization of Replacement Times Using Imperfect Monitoring Information," *IEEE Trans. on Reliability*, Vol. 52, No. 4, (December) 2003, pp. 523-533.
13. G. Heinrich, U. Jensen, "Bivariate Lifetime Distributions and Optimal Replacement," *Mathematical Methods of Operations Research*, Vol. 44, 1996, pp. 31-47.

BIOGRAPHY

Peter A. Sandborn, Ph.D.
University of Maryland
Department of Mechanical Engineering
CALCE Electronic Products and Systems Center
College Park, MD 20742

e-mail: sandborn@calce.umd.edu

Dr. Sandborn is an Associate Professor and the Research Director for the CALCE Electronic Products and Systems Center (EPSC) at the University of Maryland. His interests include technology tradeoff analysis for electronic packaging, virtual qualification of electronic systems, parts selection and management for electronic systems including electronic part obsolescence forecasting and management, and system life cycle and risk economics. Prior to joining the University of Maryland, he was a founder and Chief Technical Officer of Savantage, Inc. Dr. Sandborn has a Ph.D. degree in electrical engineering from the University of Michigan and is the author of over 80 technical publications and several books on multichip module design and electronic parts. He is an Associate Editor for the IEEE Transactions on Electronics Packaging Manufacturing.