AN APPLICATION OF DIAGNOSTIC INFERENCE MODELING TO VEHICLE HEALTH MANAGEMENT

Antony Bartolini, ARINC, (410) 266-2036, abartoli@arinc.com
Dr. John W. Sheppard, ARINC, (410) 266-2099, jsheppar@arinc.com
Thomas E. Munns, ARINC, (410) 266-2998, tmunns@arinc.com

Abstract

In this paper we discuss the approach we have taken to applying diagnostic modeling and associated reasoning techniques to the problem of diagnosing and prognosing faults as part of vehicle health management systems. We present a brief background of diagnostic fault modeling based on lessons learned from ongoing research by ARINC in cooperation with NASA Langley Research Center, as part of the NASA/FAA Aviation Safety Program. We discuss the application of these techniques and possible implementation scenarios to commercial aircraft health management. We identify information sources available on a typical commercial transport and discuss methods for evaluating them, either singly or in combination, to establish knowledge of the current or predicted health state of the aircraft.

Keywords: diagnostic inference model, system diagnosis, integrated diagnosis, diagnostic modeling, vehicle health management

1 Background

1.1 NASA/FAA Aviation Safety Program

NASA's Aviation Safety Program (AvSP) is a partnership that includes NASA, the Federal Aviation Administration (FAA), the aviation industry and the Department of Defense. AvSP was established by NASA in 1997 in response to a report from the White House Commission on Aviation Safety and Security [1]. The goal of the AvSp is "to develop and demonstrate technologies that contribute to a reduction in the aviation fatal accident rate by a factor of 5 by year 2007 and by a factor of 10 by year 2022."

The AvSP emphasizes technologies to reduce the occurrences of accidents as well as technologies to decrease injuries when accidents do occur. To achieve this end, the program includes research to pursue several ambitious goals, including:

- Reducing accidents and incidents caused by human error
- Predicting and preventing mechanical and software malfunctions

- Eliminating accidents involving hazardous weather and controlled flight into terrain

A key element of the AvSP is the Single Aircraft Accident Prevention (SAAP) initiative. The goal of SAAP is to develop safety-enabling technologies for aircraft and airborne systems in order to:

- Prevent critical system design anomalies
- Prevent critical system catastrophic failure
- Prevent loss of control in adverse flight conditions

Part of the SAAP initiative has focused on vehicle health management systems to enable failure prevention through early fault identification and real-time diagnostics.

1.2 Vehicle Health Management

Modern commercial transports are configured with sophisticated electronic, propulsion, and flight control data systems. In recent years, an increased emphasis has been placed on the potential for using these data capabilities, in conjunction with emerging sensor, data processing, and conditioning technologies for vehicle health management during flight. Implementation of such vehicle health management technologies is expected to enable operators to identify maintenance trends, anticipate component maintenance problems, implement and evaluate corrective actions, and assess performance over time. The ultimate goal is to identify and correct (or mitigate) performance or airworthiness problems before they compromise safety.

Integrated vehicle health management systems would consist of on-board systems for sensing, real-time diagnostics, and prognostics within line maintenance cycles as well as ground-based systems for longer-term diagnostics and prognostics. The onboard systems would include a variety of sensors; sensor data conditioning units; on-board diagnostic processors and algorithms; and interfaces with on-board power, data, and communications systems. The ground systems would include diagnostic and prognostic processors and algorithms, communications systems, and links to airline maintenance history records.

ARINC has undertaken research and development in vehicle health management systems to identify causal precursors to system and component failures and allow timely intervention to correct underlying malfunctions. The primary objective is to develop and demonstrate an aircraft health management system that can be integrated onboard an aircraft to provide real-time condition diagnosis and prognosis. This paper describes one key aspect of vehicle health management, the development of diagnostic models.

2 System Diagnostics

The complexity of modern systems is putting new demands on system maintenance. Every system, whether airplane, radio, or computer, has a mission to perform. The primary goal of system maintenance is to keep the system available for that mission. When the system fails, the job of maintenance is to diagnose and repair the system as rapidly as possible to return the system to correct operation. But diagnosing failures in complex systems requires analyzing characteristics of that system in great detail.

In the early 1980s, industry and government developed several initiatives to help keep pace with the growing complexity of systems and diagnostics. From these programs, groups in industry, government, and academia have developed useful testing and diagnostic tools, some of which are becoming well recognized in the test and diagnosis community. Unfortunately, each of these initiatives treated only one aspect of the life-cycle test and diagnosis problem or treated each aspect as a separate issue. None of the initiatives significantly addressed the underlying philosophy or impact of integration. Instead, integrated diagnostics was treated as nothing more than file and data sharing. While sharing data files is an important part of integrated diagnostics, software reuse and concurrent engineering are the keys to the concept.

There are three aspects of diagnosis considered in this paper, all of which apply to both on-line monitoring and off-line diagnosis.

- *Detection* refers to the ability of a diagnostic strategy to identify that a failure in some system has occurred
- *Localization* is the ability to say that a fault has been restricted to some subset of the possible causes
- *Isolation* is the identification of a specific fault through some test, combination of tests, or diagnostic strategy

The following sections detail the approach taken to develop diagnostic models for a vehicle health management application. This approach has focused on tests as information sources and the proper management of the associated information.

2.1 Diagnostic Modeling

Frequently, test engineers define a system-level diagnostic process that is independent of the design and manufacturing process. The first step, for example, is to develop BIT or built-in self test (BIST) for initial detection and localization of faults. These tests, which are embedded in the system itself, when used with other tests, could localize faults to a level sufficient to take action. Subsequent steps apply a battery of automatic and manual tests to the system (or subsystem). Eventually, these tests might identify the unit within the system or subsystem suspected of containing the fault. The unit is then tested to find the faulty subunit. Once a unit or subunit is separated from the system, maintainers frequently use specialized equipment (usually from the unit manufacturer) to test it.

Despite improvements in BIT, BIST, and automatic testing, they typically have not provided maintainers with comprehensive diagnostic procedures. Instead, they rely on part screening and special test approaches focused on individual parts or components, which adequately identify proper system function but do not adequately isolate faults when the system does not function properly. This approach provides insufficient diagnostic information to localize and isolate anomalous behavior at the system level because it fails to account for the complex interactions among system components.

In developing an alternative, we focused on ideas developed in *integrated diagnostics* programs, emphasizing the application of structured approaches to system testing and diagnosis. The three objectives of integrated diagnostics programs include:

- Maximizing reuse of design and test data, information, knowledge, and software
- Integrating support equipment and manual testing, to provide complete coverage of diagnostic requirements
- Integrating available diagnostic information, to minimize required resources and optimize performance

Our research focuses on applying a uniform method for representing and working with diagnostic information: One model type represents the system at all levels of detail. Using this model, test engineers can determine BIT requirements, define test programs for automatic test equipment, and guide the manual troubleshooting process.

The modeling approach that we use captures diagnostic inferences that arise from test information and modeling that information with respect to a set of diagnostic conclusions [3]. During troubleshooting, the information gathered from performing the series of tests is combined to make a diagnosis. Relationships between tests and conclusions results are defined in a *diagnostic inference model*. The models are hierarchical, in that a conclusion in one model can be used to invoke a lower-level model. The rules for handling each model and submodel are the same regardless of position in the hierarchy.

The diagnostic inference model [3] represents the problem to be solved via the relationship between test information and potential diagnoses. Tests provide information, and diagnostic inference combines information from multiple tests using several logical and statistical inference techniques. The structure of the diagnostic inference model facilitates the computation of testability measures and derivation of diagnostic strategies.

A diagnostic inference model has two basic elements: *tests* and *conclusions*. Tests include any source of information that can be used to determine the health of a system. Conclusions typically represent faults, including hardware fault modes, functional failures, specific nonhardware failures (such as bus timing), and specific multiple failures. A conclusion may also indicate the absence of a failure indication (no fault). Information obtained during testing might be a consequence of observing the system operation or a response to a test stimulus. Thus observable symptoms of failure are included in the diagnostic inference model as tests. Including these symptoms allows the analysis of situations involving information sources in addition to formally defined tests.

Capturing the relationships between tests and diagnoses provides a knowledge representation that can be processed by a reasoning system for health management. The type, amount, and quality of test information are considered when performing diagnosis. Initially, equal quality among test results is assumed. In other words, it is assumed that every test outcome actually reflects the state of the unit being tested. In practice, this assumption is often relaxed to allow a measure of confidence to be associated with each test.

The approach we took in developing diagnostic inference models for aircraft vehicle health management included establishing the fault universe, identifying appropriate tests, and capturing the test-to-fault relationships. Specific examples provided in the following sections refer to our recent work on health management of commercial aircraft landing gear, wheels, and brakes.

2.1.1 Establishing Fault Universe

A fundamental element of model based diagnostic systems is a well-defined fault universe, that is, an agreed upon and explicit set of diagnoses deemed worthy of the resource expenditures required to assess them. For commercial aircraft a great deal of baseline information—MSG-3 reports, minimum equipment lists, and maintenance and fault isolation manuals—is available to establish a fault universe. Some of these sources exist because of the strong focus by aircraft manufacturers and the FAA on establishing and maintaining aircraft safety, and some exists as integral parts of the aircraft design and implementation cycle.

Airlines recommend initial maintenance tasks for new aircraft based on a detailed analysis approach [4]. Each major subsystem is considered by a Maintenance Steering Group (MSG) sanctioned by the Air Transport Association (ATA). These groups consist of senior maintenance engineers from each carrier that will operate the aircraft type, as well as representatives of the manufacturer and the FAA. The groups identify significant maintenance tasks in critical systems using a rigorous evaluation process based on subsystem function, potential failure modes, and consequences of failure (e.g., affects safety, undetectable, operational impact, economic impact). As a result of the MSG-3 activities, second- and third-generation transport aircraft have formalized fault tree analyses for each major subsystem. These are invaluable in the development of health-management systems that target critical components and are consistent with the air carriers' maintenance programs.

The MSG-3 contains maintenance program development data for each major subsystem as represented by ATA chapter designations. For instance, chapter 32 includes landing gear systems. The chapter is broken down further into sections that address landing gear position and warning systems, wheels and brakes, steering, extension and retraction, and any other subsystems associated with landing gear. As an example of the next level in the hierarchy, the wheels and brakes section includes a subsection that provides a system breakdown and functional description of the hydraulic brake system. The hydraulic brake system section includes a detailed description of system functions, functional failures, the effect of failures, and the causes of failures in a tabular format. From this information table the primary function of the hydraulic brake system—to provide proper braking force—can be identified, along with potential functional failure conditions—brake failing to engage and brake failing to release. The effects of each of these failures on the hydraulic brake system are described. In the case of the brakes failing to engage the result is a reduction of braking effectiveness and in the case of brakes failing to release the result is dragging brakes. Finally the failure causes of interest (i.e., those that must be addressed by the maintenance program derived from the MSG-3 document) are listed. Brake failure, fuse set (hydraulic fuses), anti-skid valve failed open, and shuttle valve fails, are important causes of the functional failure brakes fail to engage. We ensured that the models we create strongly support the airlines' maintenance directives by capturing the failure causes as faults for our diagnostic models.

Analysis of the hierarchy of the MSG-3 report, as described in the previous paragraph, provided insight into work that has already been done to establish the system's fault universe. Although the MSG-3 reports do not provide exhaustive detail about all of the ways the system can fail, they do provide important information to start the modeling process.

Another source of diagnoses to consider in establishing the fault universe for diagnostic models is the Master Minimum Equipment List (MMEL) for the particular aircraft or aircraft subsystems being modeled. The MMEL is developed by the FAA with assistance from the manufacturer. It describes the conditions under which non-operational or faulty system components can prevent dispatching the aircraft. Conversely, the MMEL also describes the conditions and limitations for dispatching aircraft with known faulty components, depending on verified evidence that redundant systems are operational. An example, the MMEL for the Boeing 757 requires that seven of eight wheel brakes work and if one brake has failed (leaving seven) it must be carefully disconnected and capped according to the established procedure, and the landing performance must be degraded appropriately. Airlines have a refined version of the MMEL, referred to as the Minimum Equipment List (MEL), which is reflective of the configuration of their aircraft. The MEL would include MMEL items plus details that were different for different configurations of the same aircraft in the air carrier's fleet. MMEL and MEL items provide high level definitions of the health-state of the system that can be used to define diagnostic conclusions for modeling purposes.

Maintenance Manuals (MM) and Fault Isolation Manuals (FIM) also provide a great deal of information from which diagnoses can be derived. The manufacturers develop a MM, which includes a description of subsystem function, airworthiness limitations, certification maintenance requirements (CMR)[1], and servicing and lubrication requirements. Because the MSG-3 process does not continue after certification, new maintenance or modification tasks resulting from service experience (introduced through FAA Airworthiness Directives and Advisory Circulars, or manufacturer's All-Operators Letters and service bulletins), are formalized in the manufacturers' MM. Analysis of these manuals provided additional potential faults for inclusion in the fault universe.

The FIM contains troubleshooting procedures intended to lead maintenance personnel through the diagnostic process, from detection to isolation. The faults listed in the manual are obvious

---

[1] Certification maintenance requirements (CMRs) are required periodic tasks that are established during airworthiness certification as operating limitations of the type certificate.

candidates for the target fault universe. It is also possible, though with some model processing performance tradeoffs [5], to create a diagnostic inference model that captures the test strategy explicitly described by a fault tree. Thus, the fault strategies presented in the static fault trees in the FIM could be incorporated directly into diagnostic inference models. Many of the test procedures in the FIM include placing the aircraft in particular states from which to run BIT or other tests. Because these states often cannot be achieved in the air, many diagnostic models derived from ground based fault trees will only be of use for ground-based health management systems.

Other information sources that are useful in the development of diagnostic inference models include schematics, theory of operations, block diagrams, and direct interaction with maintenance personnel and engineers.

2.1.2 Tests

The number of test information sources available for analysis by the on-board component of the vehicle health management system that can be used to derive vehicle health information varies widely with aircraft type and configuration. Current commercial aircraft include a number of information sources that can be accessed via standard avionics buses. Where possible, we attempted to use sensor signals as close to the source of the sensor data as possible. For example, subsystems between the wheel-speed transducers on the landing gear and the crew display in the cockpit translate continuous values of wheel speed into integer values from 0 to 9 for display. The integer values do not provide as much useful information about the health-state of the system as the raw wheel-speed numbers. In general, for cases such as this, we have found that it is beneficial to access the data at a point prior to conversion.

As with the on-board system, the number of test information sources available for analysis by the ground-based component of the vehicle health management system depends on type and configuration of aircraft. In addition, the maintenance data logging practices of the airlines have great impact on the data available for diagnostic modeling. Along with data captured on board, we also mined historical maintenance data for failures not identified in either the MEL or MSG (these can be added to the fault universe if appropriate), failure frequencies, new ways of identifying or isolating faults as reported by maintenance personnel, and other information helpful to the diagnostic modeling process. We found that if the airline has an active Flight Operations Quality Assurance (FOQA) program additional information sources will be available for diagnostic models, although the data are not specifically captured for this purpose. Currently, total link bandwidths and the volume of other communications traffic limit the bandwidth available for communication of health management information between aircraft and ground stations. As the communication links improve and the aircraft become active nodes of a larger networked system, many of the differences between on-board and off-board information sources will decrease or disappear.

In many cases the information sources will be sufficient to detect a fault but not to isolate or localize that fault, except at a very coarse level. In these cases, additional tests could improve the isolation and localization faults. This can be accomplished either by using the existing information sources (fusing information sources to create new tests) or adding instrumentation to the original aircraft (changing the inherent information available). These added tests can be evaluated, without actually changing the aircraft, by modifying the diagnostic models to incorporate proposed test alternatives. Alternative information sources and modified diagnostic inference models are evaluated in terms of the required performance of the diagnostic system. Diagnostic models of this sort allow decisions to be made during forward-fit and retrofit implementations of the diagnostic systems.

2.1.3 Capturing Test-to-Fault Relationships

Having established a fault universe and identified tests to use to detect, localize, and isolate target faults, we then captured the test-to-fault relationships in diagnostic inference models. The diagnostic modeling representations of the IEEE standard, *Artificial Intelligence Exchange and Service Tie to All Test Environments* (AI-ESTATE), provides the needed flexibility by enabling the capture of multi-outcome tests with differing lists of faults-to-clear and faults-to-indict for each of the test outcomes [2]. In general, AI-ESTATE provides a methodology for developing diagnostic systems that will be interoperable, have transportable software, and move beyond vendor- and product-specific solutions.

Confidence levels were assigned to the test-to-fault linkages to further refine the representation of subtle test-to-fault relationships. With a sound information and test infrastructure in place, the driving issue became the identification of test-to-fault relationships for the diagnostic models. In some cases, fault simulations could be used, but in many cases human expertise was required to establish the test-to-fault relationships.

In complex systems, fault simulation could be prohibitively expensive or complex, especially for systems comprising many different technologies such as hydraulics, pneumatics, and electronics. Fault simulations require very accurate propagation of the resulting behavior once a fault has been injected. In complex systems of mixed technology, simulations often do not include the detail required to propagate injected faults accurately. If the resolution required of the simulation is too high, available computational resources could be inadequate to establish the desired test-to-fault relationships. Finally, human experts often have intuitive simulations of fairly complex relationships based on experience and observation. This understanding of system behavior, if captured, can be useful in developing accurate diagnostic models. We used a combination of human expertise, along with computer fault simulation to derive test-to-fault relationships.

As described in a previous section, the FIM often contains static fault trees for aircraft troubleshooting by maintenance personnel while the aircraft is on the ground. Although a health management system should detect and if possible localize and isolate many of the same faults targeted by the FIM, the reasoning process may not be of a static nature. Therefore, the test-to-fault relationships expressed in the fault trees could be included in the diagnostic models, but the predefined decision trees might not be explicitly incorporated into the processing algorithms of the diagnostic reasoners. Many FIM procedures cannot be safely performed on-board a flying aircraft. The diagnostic knowledge implicit in these procedures may be useful on the ground, but inappropriate in the air. In these cases, we performed detailed diagnostic analyses with cross checking by domain experts.

This section has shown that accurate and complete diagnostic inference models, including an established fault universe, appropriate tests, and thorough understanding of test-to-fault relationships are central to the development of a vehicle health management system. The creation of accurate diagnostic models is labor intensive, but with an appropriate diagnostic reasoner, the benefits to reduced system life cycle maintenance costs and improved safety often make the effort worthwhile [6].

2.2 Accumulation of Evidence

The algorithm applied for drawing inferences from actual test information [7][8] uses a modification of the Dempster-Shafer [7][9] statistical inference, which is derived from Bayesian inference theory. The Dempster-Shafer modifications included (1) limiting the fault universe to the set of simple conclusions (i.e., diagnoses) in the model and (2) defining a special conclusion, *unanticipated result* [7].

The two extremes of a credibility interval, called *support* and *plausibility*, are calculated for every conclusion. The probability that a given conclusion is true lies between its support and plausibility values. A test outcome *supports* a conclusion (and thereby the associated fault) when the outcome indicates the detection of the fault associated with that conclusion. A test outcome *denies* a conclusion if it eliminates the conclusion from consideration (denial is the complement of plausibility). Confidence values are assigned to test outcomes to compute support and denial values.

Because the support value depends strongly on previously normalized data, the Dempster-Shafer calculations exhibit a *temporal-recency effect*, that is, recent events have a greater impact on the evidential calculation than distant events. Because of this undesirable property, we explored alternative approaches to reasoning under uncertainty. We wanted to be able to base our inferences on the information flow model, assign confidences to test outcomes, and perform consistent inference, independent of temporal ordering. Specifically, we derived a simplified approach to reasoning with uncertain test data and discovered that we had re-derived a relatively old method—*certainty factors*. Our application of certainty factors to system diagnosis differs from Dempster-Shafer by assigning the full confidence value to all conclusions either supported or denied rather than apportioning confidence to the supported conclusions. Obviously, support is applied to a conclusion only if the test outcome actually supported that conclusion, and denial is applied only if the test outcome actually denied the conclusion.

Updating support and denial over time is also different from the Dempster-Shafer approach and is straightforward; it is similar to combining probabilities. We updated these measures by adding the current support or denial to the previously accumulated support or denial and subtracting the product of the two. We determined certainty in a conclusion by subtracting the accumulated denial from the accumulated support. We then rescaled the resulting certainty value between zero and one, so the value could be interpreted like a probability.

The primary advantages to using certainty factors rather than Dempster-Shafer include reduced computational complexity and sequence independence in determining support and denial for each of the conclusions. Dempster-Shafer's primary advantage is a firmer grounding in probability theory and a larger base of practical experience demonstrating acceptable behavior.

2.3 Prognosis

The approach we have taken to performing system prognosis is an extension of the model-based diagnostic approach. Specifically, we began by observing that prognosis and diagnosis are two aspects of the same problem. Diagnosis consists of determining the current health state of a system given a set of test results. Prognosis also consists of determining the health state of a system; however, this health state is expected to occur sometime in the future.

Most current approaches to prognosis apply one of two techniques. The first is reliability based, in which the reasoner considers the failure rate of elements and components within a system, ties those failure rates to operational data, and projects how much time remains until a failure is likely to occur. Little to no test information is used in the reliability-based approach to refine the projection. The second approach relies on detailed physics-based models of the system and processes these models from a current known state. This "physics-of-failure" approach is generally considered to be highly accurate (given an accurate physical model); however, the computational resources required to process these models for any system of reasonable size are substantial.

The approach to health management and diagnosis that we have taken is an "information-centric" approach. In other words, we have focused on extracting information from test results and relating that information to the potential diagnostic conclusions (e.g., faults) in the model. Our approach to prognosis is also "information-centric." Specifically, we have focused on projecting test results rather than projecting the occurrence of a fault. By projecting changes in test results over time, we have been able to abstract the prediction problem from the diagnostic problem and apply an alternative that combines the low cost of reliability-based prognostics with the accuracy of physics-based prognostics.

3 Conclusions

ARINC has undertaken research and development to develop and demonstrate an aircraft vehicle health management system that can be integrated onboard an aircraft to provide real-time condition diagnosis and prognosis. The approach to health management and diagnosis that we have taken is "information-centric," focused on extracting information from test results and relating that information to potential diagnostic conclusions. We have already achieved significant advances in integrated diagnostics systems by focusing on tests as information sources and properly managing the associated information.

By using problem encapsulation, defining interface boundaries, developing exchange formats and specifying standard services, AI-ESTATE provides a methodology for developing diagnostic systems that will be interoperable, have transportable software, and move beyond vendor- and product-specific solutions.

Emphasis has been placed on the application of a uniform method for representing and working with diagnostic information (i.e., one model type represents the system at all levels of detail). Accurate and complete diagnostic inference models, including an established fault universe, appropriate tests, and thorough understanding of test-to-fault relationships were the key to the development of a vehicle health management system.

The algorithm applied for drawing inferences from actual test information uses a modification of the Dempster-Shafer statistical inference, modified by limiting the conclusion space to the set of simple conclusions (i.e., diagnoses) in the model and defining a special conclusion, *unanticipated result.* Other methods to handle uncertainty were investigated because the Dempster-Shafer calculations exhibit a temporal-recency effect (i.e., more recent events have a greater impact on the evidential calculation than more distant events). Certainty factors were applied to base inferences on the information flow model; assign confidences to test outcomes; and perform consistent inference, independent of temporal ordering.

The approach to performing system prognosis is an extension of the model-based diagnostic approach, that is, prognosis and diagnosis are two aspects of the same problem. Diagnosis consists of determining the current health state of a system given a set of test results. Prognosis also consists of determining the health state of a system; however, this health state is expected to occur sometime in the future.

These developments have allowed us to develop a prototype aircraft vehicle health management system based on an integrated diagnostic approach.

4 References

[1] White House Commission on Aviation Safety and Security. 1997. *Final Report to the President.* Washington, D.C.: U.S. Government Printing Office. Available on-line at http://www.fas.org/irp/threat/212fin~1.html /

[2] IEEE P1232/D4.0. 2000. *Draft Standard for Artificial Intelligence Exchange and Service Tie to All Test Environments (AI-ESTATE)*, Piscataway, NJ: IEEE Standards Press.

[3] Simpson, W. and Sheppard, J. 1994. *System Test and Diagnosis,* Boston, MA: Kluwer Academic Publishers.

[4] *Airline/Manufacturer Maintenance Program Development Document (Revision 2).* 1993 Maintenance Steering Group 3 (MSG-3) Task Force, Washington, DC: Air Transport Association of America.

[5] Bearse, T. M. 1998. Deriving a Diagnostic Inference Model from a Test Strategy. Chapter 4 in *Research Perspectives and Case Studies in System Test and Diagnosis,* J. W. Sheppard and W. R. Simpson, eds. Boston, Mass.: Kluwer Academic Publishers.

[6] Kent, R. M. and D. A. Murphy. 2000. *Health Monitoring System Technology Assessments: Cost Benefits Analysis*, NASA/CR-2000-209848, Hampton, VA: NASA Langley Research Center.

[7] Dempster, A. 1968. A generalization of Bayesian inference. *Journal of the Royal Statistical Society*, Series B, pp. 205–247.

[8] Sheppard, J. 1994. Maintaining Diagnostic Truth with Information Flow Models. Pp 447–454 in *AUTOTESTCON '94 Conference Record*, New York: IEEE Press.

[9] Shafer, G. 1976. *A Mathematical Theory of Evidence*, Princeton, NJ: Princeton University Press.