

Detection and Estimation of Multiple Fault Profiles Using Generalized Likelihood Ratio Tests: A Case Study

Joshua D. Carl* Ashraf Tantawy* Gautam Biswas*
Xenofon D. Koutsoukos*

* *Institute for Software Integrated Systems, Vanderbilt University,
Nashville, TN 37235 {carljd1, tantawam, biswas,
koutsord}@isis.vanderbilt.edu*

Abstract: Aircraft and spacecraft electrical power distribution systems are critical to overall system operation, but these systems may experience faults. Early fault detection makes it easier for system operators to respond and avoid catastrophic failures. This paper discusses a fault detection scheme based on a tunable generalized likelihood algorithm. We discuss the detector algorithm, and then demonstrate its performance on test data generated from a spacecraft power distribution testbed at NASA Ames. Our results show high detection accuracy and low false alarm rates.

Keywords: Fault Detection, Fault Diagnosis, Fault Identification, Fault Location

1. INTRODUCTION

Faults and degradation in dynamic systems can be attributed to different fault profiles that have different temporal characteristics: (1) abrupt persistent faults, (2) abrupt intermittent faults, and (3) incipient faults. Two interesting challenges that we discuss are: (1) how to extend detection schemes to include abrupt intermittent faults, and (2) how to combine abrupt, incipient, and intermittent fault detection into an integrated on-line detection scheme.

This paper develops a general framework that enables on-line detection of all three fault profiles using an integrated detector. The detector employs change detection theories that apply to systems with stochastic behaviors attributed to uncertainty in the model parameters and measurement noise. In our work, we assume measurement noise can be modeled by Gaussian distributions with zero mean. We use these assumptions to obtain closed form expressions for the estimators and a recursive expression for the fault detector. We study the performance of this detector on the ADAPT power distribution system at NASA Ames.

Section 2 of this paper presents a brief overview to the case study based on the ADAPT testbed at NASA Ames. Section 3 presents the conceptual framework for the detection problem, section 4 describes our approach and the algorithms we have implemented for online detection of faults, and section 5 presents the results of our case study. The final section presents our conclusions.

2. CASE STUDY OVERVIEW

We perform a case study of the fault detector on data generated from the NASA ADAPT-Lite Electrical Power System (EPS), as part of the DXC'10 diagnosis competition. A complete description of the DXC'10 competition

can be found in [Kurtoglu et al., 2010] and the algorithm evaluation metrics in [Kurtoglu et al., 2008].

The EPS supplies power to spacecraft systems and payloads. The EPS schematic in Figure 1 shows a battery connected to a load bank through a set of switches, circuit breakers and an inverter. Since the dynamics of the inverter (a fast switching system that converts DC voltage to AC) was not a factor in this competition, the rest of the system represents a power source that is connected to resistive loads that can be reconfigured using the switches. Therefore, the system behavior is static, and a fault in any of the system components produces an identical profile in the related sensor data.

The competition was sponsored by researchers at NASA Ames, and was designed to mimic a live situation so all detection algorithms were required to support on-line analysis. Each detection algorithm was evaluated using metrics such as the fault detection time, accuracy, and false alarm rate. A fast and accurate diagnosis and fault isolation was a requirement to do well in the competition.

3. PROBLEM FORMULATION

3.1 Fault Detection Architecture




Fig. 2. Fault detection system.

Our observer-based approach for fault detection is illustrated in Figure 2. The physical system being monitored and the observer receive the same input signals, and the

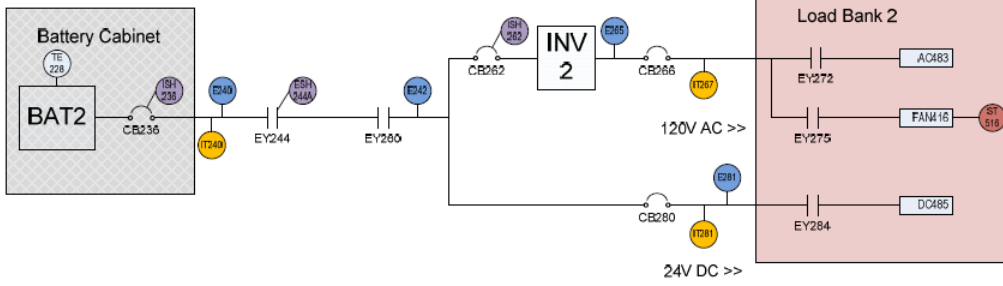


Fig. 1. NASA Electrical Power System Schematic Diagram.

system output, i.e., the actual system measurements, are labeled as $y[n]$, and the observer estimates are labeled as $\hat{y}[n]$. The system residual vector is computed as $y[n] - \hat{y}[n] = r[n]$ at timestep n . The fault detector uses hypothesis testing methods to determine if the computed residual signals imply a fault in the system. The fault detector has to be robust to measurement noise, system disturbances as well as model inaccuracies. The output of the fault detector is a vector of binary variables, b , representing the fault signature for the system, and a set of parameters, θ , that describe the change in the residual signal. A non-zero value for b_i implies that measurement i is deviant from its nominal value. Since the relation between a fault and corresponding measurement values are algebraic, the detector also provides an estimate for the residual parameters relevant to the fault type. The output of the detector is given to fault isolation and fault parameter identification units (not shown) for the completion of the fault processing. The design of the fault isolator and fault parameter identifier are discussed in [Carl et al., 2012].

In this paper we focus on the design of the fault detector. The detector is required to accomplish the following tasks: (1) decide if there is a change in the nominal behavior of the system, (2) declare whether the existing fault is an abrupt persistent, abrupt intermittent, or incipient, and (3) estimate the relevant residual parameters.

3.2 Fault Hypothesis

We assume that the signals generated by the physical system have added independent and identically distributed Gaussian Noise, represented as $w[n]$, with zero mean and an unknown variance. The variance can be calculated with knowledge of nominal system behavior and the sensor measurements before a fault occurs in the system.

When a fault occurs in the system starting at t_{inj} , the system measurements can be defined as:

$$y[n] = \begin{cases} s[n] + w[n] & t < t_{inj} \\ s_{f_i}[n] + w_{f_i}[n] & t \geq t_{inj} \end{cases} \quad (1)$$

Where $s[n]$, the nominal signal value at time step n , is known from available system behavior data, or is estimated using an observer scheme [Basseville and Nikiforov, 1993] and $w[n]$ represents the noise in the measurement that is typically attributed to the sensor. After the fault occurrence, the signal value is linked to faulty system behavior and is expressed as $s_{f_i}[n]$ for $n \geq t_{inj}$, with a corresponding noise component that is given by $w_{f_i}[n]$. In our work, we assume that the measurement noise is unaffected by

system faults, therefore $w_{f_i}[n] = w[n]$. Taking this into account, the detection problem can be expressed as:

$$\begin{aligned} \mathcal{H}_0 &: r[n] = w[n] \\ \mathcal{H}_i &: r[n] = \Delta s_{f_i}[n] + w[n] \quad i = 1, 2, \dots, m \end{aligned} \quad (2)$$

where \mathcal{H}_0 is the null hypothesis of no fault, \mathcal{H}_i is the alternative (fault) hypothesis, m is the number of faults, and $\Delta s_{f_i}[n] = s_{f_i}[n] - s[n]$ represents the deviation in the measurement as a result of the fault.

We formulate the detection problem for three different fault profiles. The detector needs to estimate a variety of parameters for each fault type, and the detection problem for each fault is defined by the fault profile and the set of parameters associated with the profile.

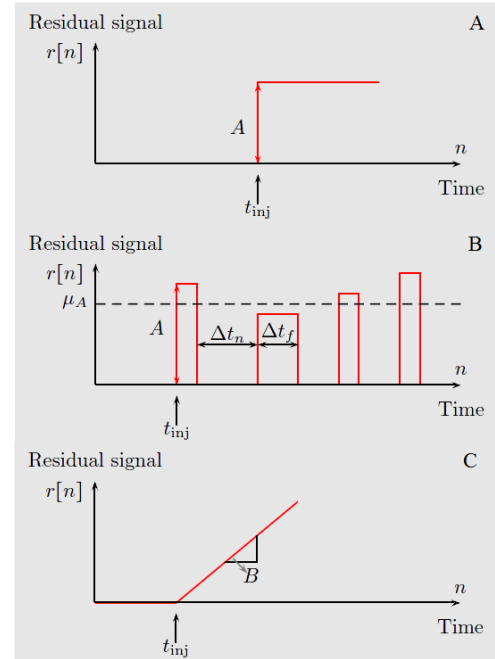


Fig. 3. Idealized fault profiles. Noise is removed from the residual signal for clarity.

3.3 Fault Profiles

Abrupt Persistent The abrupt persistent fault profile, shown in Figure 3A, is characterized by the nominal signal changing by an unknown positive or negative additive fixed value. For an abrupt persistent fault the detector needs to estimate the fault time of injection, t_{inj} , and the change

in the magnitude of the signal, A , caused by the abrupt fault. The residual signal for the fault is expressed as:

$$r[n] = A + w[n]. \quad (3)$$

Abrupt Intermittent An abrupt intermittent fault profile, shown in Figure 3B, is modeled as a repeated abrupt persistent fault that resets itself after a random time interval. The fault persistence time, Δt_{fi} , and the inter-arrival time, Δt_{ni} , for each fault repetition are drawn from exponential distributions of $\exp(\mu_f, t_f)$ and $\exp(\mu_n, t_n)$, respectively. The change in residual signal magnitude, A , caused by the fault is drawn from a Gaussian distribution with mean μ_A and variance σ_A^2 . The detector needs to estimate the fault time of injection, t_{inj} , the mean residual signal magnitude, μ_A , mean persistence time for the fault, μ_f , and mean inter-arrival time for the fault, μ_n .

The residual signal for the fault can be expressed as:

$$r[n] = AZ[n] + w[n], \quad (4)$$

where the function $Z[n]$ is a binary random process representing the presence or absence of the fault, defined by:

$$Z[n] = \begin{cases} 0 & \text{fault absent} \\ 1 & \text{fault present} \end{cases}. \quad (5)$$

Incipient An incipient fault profile, shown in Figure 3C, is a linear change (positive or negative) in the sensor signal. Incipient faults can be approximated by a linear profile, because they evolve slowly in time. For an incipient fault the detector needs to estimate the time of injection, t_{inj} , and the slope of the signal, M . The residual signal for the fault can be expressed as:

$$r[n] = Bn + w[n] \quad (6)$$

where $B = MT_s$, M is a constant representing the slope of the drift, and T_s is the sampling period.

4. TECHNICAL APPROACH

The detector architecture is shown in Figure 4. The detector receives inputs of a residual data set, r . It detects faults using a Generalized Likelihood Ratio (GLR) test statistic, where the unknown parameters are replaced by their respective Maximum Likelihood (ML) estimates. If there is a fault the detector outputs the change in the residual magnitude, ΔA , and the fault injection time, t_{inj} , to the data vector catalog. The catalog tracks the individual inputs over time and stores them as vectors. The fault profiler uses the vector of residual magnitudes, $\hat{\mathbf{A}}$, to determine the type of fault. The type of fault, residual magnitude vector, and the vector of fault times, $\hat{\mathbf{T}}$, are given to the residual parameter estimator, which estimates the relevant parameters for the fault, θ_1 . The parameters and the change of the residual magnitude are passed out of the detector.

It is important to note a key characteristic of our detector implementation. After a fault is detected, the detector resets and uses the estimated residual signal magnitude, \hat{A} , as its new baseline and starts sampling from this reset point, ignoring the residual data it was given previously. This allows each fault type to generate its own fault profile, which is used to tell the different fault types apart.

4.1 Detector Derivation

The majority of this derivation is taken from [Basseville and Nikiforov, 1993]. We designate the Log Likelihood Ratio (LLR) for observations r from time j up to time k by:

$$S_j^k(\theta_1) = \sum_{i=j}^k \ln \frac{p_{\theta_1}(r[i])}{p_{\theta_0}(r[i])} \quad (7)$$

where θ_0 and θ_1 are the set of parameters that characterize the distributions of observations, p_{θ_0} and p_{θ_1} , before and after the change, respectively, and $r[i]$ is the signal residual. p_{θ_0} is assumed known from the nominal system behavior data and is a Gaussian distribution with zero mean and variance σ_n^2 , while p_{θ_1} is assumed unknown. We know σ_n^2 from the nominal system data. In addition, the change time is unknown to the detector. These parameters are substituted by their ML estimates, and the test statistic is given by:

$$g_k = \max_{1 \leq j \leq k} \sup_{\theta_1} S_j^k(\theta_1). \quad (8)$$

The detection time, t_a , is the minimum value of k at which $g_k > h$, where h is the detector threshold.

$$t_a = \min \left\{ k \mid \max_{1 \leq j \leq k} S_j^k \geq h \right\} \quad (9)$$

The conditional ML estimate for the change time is the value of j at which the maximum value of g_k is reached. Therefore, the conditional ML estimate for the change magnitude and time are given by:

$$(\hat{t}_{inj}, \hat{\theta}_1) = \arg \max_{1 \leq j \leq t_a} \sup_{\theta_1} \sum_{i=j}^{t_a} \ln \frac{p_{\theta_1}(r[i])}{p_{\theta_0}(r[i])} \quad (10)$$

where \hat{t}_0 is the estimated fault injection time.

Abrupt Persistent Fault Detection After a fault both the change magnitude, $\theta_1 = \hat{A}$, and the change time are unknown. After simplifications, the LLR can be written as:

$$S_j^k = \frac{\hat{A}}{\sigma_n^2} \sum_{i=j}^k \left(r[i] - \frac{\hat{A}}{2} \right) \quad (11)$$

and the test statistic is:

$$g_k = \frac{1}{2\sigma_n^2} \max_{1 \leq j \leq k} \frac{1}{k-j+1} \left[\sum_{i=j}^k r[i] \right]^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\geq}} \gamma \quad (12)$$

where γ is the abrupt fault detection threshold and \mathcal{H}_1 and \mathcal{H}_0 are the two fault hypothesis from (2).

Since we reset the GLRT detector after each fault is declared, after the first fault, assuming no fault alarms, no more faults will be detected. This type of fault will have a single element in its detection profile:

$$\text{sgn}(\hat{A}) = \{1\} \quad (13)$$

where sgn is the sign function.

Abrupt Intermittent Fault Detection In an intermittent fault multiple abrupt fault instances will be detected, so the detection problem and the test statistic are the same as the abrupt persistent fault case. Since both A for the abrupt persistent fault, and μ_A for the abrupt intermittent




Fig. 4. Fault detector architecture.

fault are unknown, there is no way to differentiate between the two types of faults using only the GLR test. We are able to discriminate between the two fault types because the detector resets its baseline measurement to the new faulty residual magnitude after each fault. Therefore, each two consecutive faults will have opposite signs for \hat{A} . We can summarize the detection profile by $\text{sgn}(\hat{A})$. The fault profile, assuming positive residual change magnitude and no false alarms is given by:

$$\text{sgn}(\hat{A}) = \{1 \ -1 \ 1 \ -1 \ 1 \ \dots\}. \quad (14)$$

A threshold count can be set for how many fault instances are required before declaring a fault to be intermittent.

Incipient Fault Detection Since the incipient fault is characterized by a continuous increase in the signal magnitude, the application of the LRT for abrupt faults will detect consecutive changes in one direction only. This is due to the reset of the LRT after it finds a fault; it will reset repeatedly as the signal continues to increase. Therefore, the detection profile for incipient faults, assuming positive residual change magnitude, is given by:

$$\text{sgn}(\hat{A}) = \{1 \ 1 \ 1 \ 1 \ \dots\}. \quad (15)$$

Similarly, a threshold can be set for how many fault instances are required before declaring a fault to be incipient. Our detection approach will also work on a non-linear signal, except that the signal model will contain more unknowns than the simple slope presented here.

4.2 Residual Parameter Estimation

In this section we present expressions for residual parameter estimators for each fault type. The results are drawn directly from classical estimation theory techniques [Kay, 1993].

Abrupt Persistent Faults The detection algorithm returns the fault injection time, t_{inj} , and the residual signal magnitude, \hat{A} , which are the only parameters to be estimated for this fault type.

Abrupt Intermittent Faults The parameters to be estimated are:

- (1) **Fault Injection Time.** The fault detection algorithm returns a vector of all the fault injection times. In the case of intermittent faults the fault injection time is the time instant at which the first fault takes place:

$$\hat{t}_{inj} = \hat{\mathbf{T}}[1]. \quad (16)$$

- (2) **Mean Residual Signal Magnitude.** The residual signal magnitude is drawn from a Gaussian distribution. Therefore, the ML estimator for its mean

$\hat{\mu}_A$ is just the arithmetic mean of the residual signal magnitude vector $\hat{\mathbf{A}}$:

$$\hat{\mu}_A = \bar{\mathbf{A}} = \frac{1}{\text{size}(\hat{\mathbf{A}})} \sum_{i=1}^{\text{size}(\hat{\mathbf{A}})} \hat{\mathbf{A}}[i] \quad i = 1, 3, 5, \dots \quad (17)$$

- (3) **Mean Time Between Faults.** The inter-arrival times can be calculated from the vector \mathbf{T} as:

$$\hat{\mathbf{T}}_n = \hat{\mathbf{T}}[i+1] - \hat{\mathbf{T}}[i] \quad i = 2, 4, 6, \dots \quad (18)$$

The inter-arrival time has an exponential distribution. Therefore, the ML estimator for its mean is just the arithmetic mean of $\hat{\mathbf{T}}_n$:

$$\hat{\mu}_n = \bar{\mathbf{T}}_n = \frac{1}{\text{size}(\hat{\mathbf{T}}_n)} \sum_{i=1}^{\text{size}(\hat{\mathbf{T}}_n)} \hat{\mathbf{T}}_n[i]. \quad (19)$$

- (4) **Mean Fault Duration.** Similarly, the fault durations are calculated as follows:

$$\hat{\mathbf{T}}_f = \hat{\mathbf{T}}[i+1] - \hat{\mathbf{T}}[i] \quad i = 1, 3, 5, \dots \quad (20)$$

and the ML estimator is given by:

$$\hat{\mu}_f = \bar{\mathbf{T}}_f = \frac{1}{\text{size}(\hat{\mathbf{T}}_f)} \sum_{i=1}^{\text{size}(\hat{\mathbf{T}}_f)} \hat{\mathbf{T}}_f[i]. \quad (21)$$

Incipient Faults The parameters to be estimated are:

- (1) **Fault Injection Time.** Similar to the case of an abrupt intermittent fault, we define the fault injection time as the time instant at which the first fault takes place:

$$\hat{t}_{inj} = \hat{\mathbf{T}}[1]. \quad (22)$$

- (2) **Drift Slope.** We note from (6) that the observations represent a linear model in the unknown parameter B . To show that, we write the vector form of (6):

$$\mathbf{r} = \mathbf{H}\mathbf{B} + \mathbf{w} \quad (23)$$

where $\mathbf{H} = [0 \ 1 \ 2 \ \dots \ N]$. We note that sample 0 corresponds to the time instant t_{inj} and sample N corresponds to the time instant NT_s , where T_s is the sampling period and N is the total number of samples used in the estimation process. The solution of the estimation problem for the linear model in (23) results in the Minimum Variance Unbiased Estimator (MVUE):

$$\hat{B} = \left(\mathbf{H}^T \mathbf{H} \right)^{-1} \mathbf{H}^T \mathbf{r} \quad (24)$$

using the expression for \mathbf{H} we obtain:

$$\hat{B} = \frac{\sum_{n=0}^{N-1} nr[n]}{\sum_{n=0}^{N-1} n^2}. \quad (25)$$

4.3 Detector Enhancements

The incipient fault detection scheme presented was based on a heuristic approach of repetitive change detection.

This process is suboptimal, since it assumes the data model for the abrupt fault. A more accurate method is to run two binary hypothesis tests in parallel, one for each fault type (abrupt and incipient). In the likely event that both detectors will fire a detection event, an additional LRT is performed between the data models of the abrupt and incipient faults to decide which fault is the most likely one. The data samples that are used in this LRT are the ones after the latest fault is declared, and the length of the data samples used is to be decided based on the required accuracy. This enhanced detection scheme is depicted in Figure 5. The abrupt fault detector has the same design




Fig. 5. Enhanced detector design for efficient detection of incipient faults.

as the one presented above and we need to design the incipient fault detector and the fault selector. We start with the incipient fault detector, where we have the binary composite hypothesis testing problem:

$$\begin{aligned} \mathcal{H}_0 : r[n] &= w[n] \\ \mathcal{H}_1 : r[n] &= Bn + w[n]. \end{aligned} \quad (26)$$

We calculate the LLR as above, noting that the ML estimator for B is given by (25). It is straightforward to show that the test statistic is given by:

$$g_k = \frac{1}{2\sigma_n^2} \max_{1 \leq j \leq k} \frac{1}{\sum_{i=j}^k i^2} \left[\sum_{i=j}^k ir[i] \right]^2 \begin{matrix} \mathcal{H}_1 \\ \geq \\ \mathcal{H}_0 \end{matrix} \gamma_c \quad (27)$$

where γ_c is the incipient fault detection threshold, and \mathcal{H}_1 and \mathcal{H}_0 are the fault hypothesis from (26).

When the fault selector receives an alarm from one of the fault detectors, it carries out an LRT between the two fault distributions:

$$\begin{aligned} \mathcal{H}_0 : r[n] &= A + w[n] \\ \mathcal{H}_1 : r[n] &= Bn + w[n] \end{aligned} \quad (28)$$

The LRT produces the following test statistic, after substituting for the ML estimators for A and B :

$$g_k = \frac{1}{2\sigma_n^2} \left[\frac{\left(\sum_{n=0}^{N-1} nr[n] \right)^2}{\sum_{n=0}^{N-1} n^2} - \frac{\left(\sum_{n=0}^{N-1} r[n] \right)^2}{N} \right] \begin{matrix} \mathcal{H}_1 \\ \geq \\ \mathcal{H}_0 \end{matrix} \gamma_{ic} \quad (29)$$

where γ_{ic} is the decision threshold between the two fault detectors, and \mathcal{H}_1 and \mathcal{H}_0 are the fault hypothesis from (28).

Equations (12), (27), and (29) make up the different components in Figure 5, where (12) is the abrupt fault detector, (27) is the incipient fault detector, and (29) is the fault selector. The abrupt fault detector and the

incipient fault detector run in parallel, each comparing one faulty case to the null hypothesis (no fault case). The fault selector is only used when both detectors detect a fault, and is used to determine which fault scenario is the most likely. An adjustable delay τ , may be introduced into the fault selector, to take into account the probability that one detector produces a fault before the other one. When both detectors find a fault, the fault injection time is considered to be the maximum of the two fault injection times. Also, the number of samples, N , is adjustable, based on the required probability of detection for a given probability of false alarm rate.

5. EXPERIMENTAL RESULTS

The DXC'10 competition has 154 fault scenarios. In each scenario, a single fault in one of the system components is injected. The fault types include abrupt, abrupt intermittent, drift. There are pictures of faults below. The incipient fault is the most challenging to detect quickly, since drift slope is very small compared to the noise variance. Table 1 shows the performance of the detector with the competition scenarios.




Fig. 6. IT281 Abrupt fault. $\hat{A} = -0.05$.

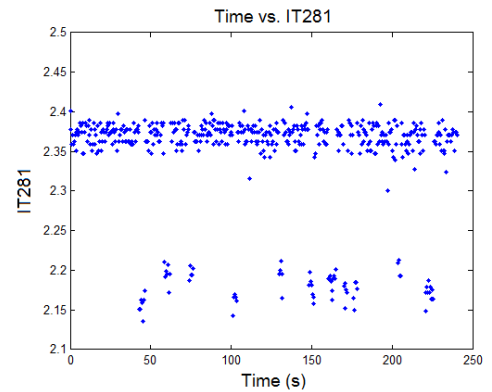


Fig. 7. IT281 Abrupt intermittent fault. $\hat{\mu}_A = -0.19$, $\hat{\mu}_f = 3.21$, and $\hat{\mu}_n = 16.14$.

The detector performed reasonably well for abrupt persistent faults. Most of the missed detections in this case were because of improper tuning of the detector. In the fault scenarios presented, different noise levels were associated with the same sensor in different scenarios. Since the detector threshold was fixed, based on the training data

Fault Type	Total Scenarios	Detected	Undetected	P_D	Remedy
Abrupt Persistent	37	29	8	0.783	Adaptive threshold
Abrupt Intermittent	35	25	10	0.714	Low γ_i
Incipient	37	19	18	0.513	LRT
Incipient with Incipient Detector	37	34	3	0.919	None required

Table 1. Detector performance, NASA DXC'10 competition.




Fig. 8. IT281 Incipient fault. $M = 0.001$.

set, the detector was not able to cope with the change in variance from one fault scenario to the other. The solution of the problem is to use an adaptive threshold, where the algorithm automatically sets the detector threshold based on the estimated variance value for the incoming data set.

The detector performance for abrupt intermittent faults was worse than for abrupt persistent faults. The performance for intermittent faults cannot be better than persistent faults since missing an abrupt fault leads also to an intermittent fault miss. The excess performance degradation is because the threshold value was set higher than necessary. That was mainly to accommodate for false alarms from the change detector. This problem could be addressed by lowering the threshold and reducing the false alarm rate simultaneously. The false alarm rate can be reduced by proper detector tuning.

We note that the detector performed poorly for the incipient faults, because we relied on the heuristic approach of multiple change detections in one direction as an incipient fault profile. Unfortunately, most of the fault scenarios in the competition for incipient faults had very small drift slope and a limited data set, where the data set was not long enough for the detector to detect enough changes to declare an incipient fault. For this reason, most of the undetected incipient faults were reported as abrupt persistent. When the incipient detector and fault selector described in (27) and (29) were added to the overall fault detection model the results for detecting incipient faults improved dramatically, as shown in the last row of Table 1.

6. CONCLUSION

Statistical models are powerful in designing fault detectors for physical systems, provided that a data set is available for the nominal system behavior and the set of faults of interest. We presented a general fault detection algorithm, based on change detection theory, which is capable of detecting abrupt, intermittent, and incipient faults.

Several enhancements are possible for the presented algorithm. An adaptive threshold that changes with the noise variance level could increase the probability of detection. The distinction between abrupt persistent and intermittent faults could be enhanced by reducing the false alarm rate of the change detector. Recursive detection and estimation statistics are also important in practical implementations to speed up the decision process, allowing for early fault declaration and parameter estimation.

ACKNOWLEDGEMENTS

We would like to thank NASA Ames for sponsoring the 2010 Diagnostic Competition.

REFERENCES

- Michele Basseville and Igor V. Nikiforov. *Detection of Abrupt Changes: Theory and Applications*. Prentice Hall, Inc., 1993.
- Joshua D. Carl, Daniel L.C. Mack, Ashraf Tantawy, Gautam Biswas, and Xenofon D. Koutsoukos. Fault detection and isolation for spacecraft systems: An application to a power distribution testbed. In *8th IFAC SAFE-PROCESS*, 2012.
- Steven M. Kay. *Fundamentals of Statistical Signal Processing, Volume 1: Estimation Theory*. Prentice Hall Signal Processing Series. Prentice Hall PTR, 1993.
- Tolga Kurtoglu, Sriram Narasimhan, Scott Poll, David Garcia, Lukas Kuhn, Johan de Kleer, and Alexander Feldman. The Diagnostic Challenge Competition - DCC'09. Technical report, NASA Ames Research Center, Palo Alto Research Center, Delft University of Technology, 2008. <https://c3.nasa.gov/dashlink/projects/36/>.
- Tolga Kurtoglu, Sriram Narasimhan, Scott Poll, David Garcia, Lukas Kuhn, Johan de Kleer, and Alexander Feldman. Second International Diagnostic Competition (DXC'10), Industrial Track, Diagnostic Problem Descriptions. Technical report, NASA Ames Research Center, 2010. <https://www.phmsociety.org/competition/dxc/10>.