

DETECTION LIMITS FOR LINEAR NON-GAUSSIAN STATE-SPACE MODELS

Gustaf Hendeby * Fredrik Gustafsson *

* *Division of Automatic Control
Department of Electrical Engineering,
Linköpings universitet, SE-581 83 Linköping, Sweden
{hendeby, fredrik}@isy.liu.se*

Abstract: The performance of nonlinear fault detection schemes is hard to decide objectively, so Monte Carlo simulations are often used to get a subjective measure and relative performance for comparing different algorithms. There is a strong need for a constructive way of computing an analytical performance bound, similar to the *Cramér-Rao lower bound* for estimation. This paper provides such a result for linear non-Gaussian systems. It is first shown how a batch of data from a linear state-space model with additive faults and non-Gaussian noise can be transformed to a residual described by a general linear non-Gaussian model. This also involves a parametric description of incipient faults. The *generalized likelihood ratio* test is then used as the asymptotic performance bound. The test statistic itself may be impossible to compute without resorting to numerical algorithms, but the detection performance scales analytically with a constant that depends only on the distribution of the noise. It is described how to compute this constant, and a simulation study illustrates the results. Copyright © 2006 IFAC

Keywords: Detector performance; Fault detection; Linear Systems; Non-Gaussian process; Performance analysis; State-space models; Stochastic fault detection

1. INTRODUCTION

This paper studies fault detection in linear non-Gaussian systems. It is first shown how a batch of data from a linear state-space model with additive faults and non-Gaussian noise can be transformed into a residual described by a general linear non-Gaussian model of the form

$$\mathbb{R}_t = \bar{H}_t^\theta \theta + \bar{H}_t^v \mathbb{V}_t. \quad (1)$$

This transformation is based on either prior knowledge, estimation based on past data, or a parity space approach. Here, \mathbb{V}_t captures the stochastic effects in the system possibly with coloring \bar{H}_t^v . Further, θ is a parameter vector in a smooth parameterization of incipient fault profiles, which affects the system through the matrix \bar{H}_t^θ (see Sec. 2). The generating system is said to be fault free if $\theta = 0$ and otherwise faulty,

i.e., fault detection thus turns into a hypothesis test, (Kay, 1998; Basseville and Nikiforov, 1993),

$$\begin{cases} \mathcal{H}_0 : \theta = 0, \\ \mathcal{H}_1 : \theta \neq 0. \end{cases} \quad (2)$$

This test may be repeated for all possible faults, by using different \bar{H}_t^θ , for the purpose of diagnosis.

The goal of this paper is to compute an asymptotic upper bound for the detection probability P_D for a given false alarm rate P_{FA} ,

$$P_D(\hat{d}(\mathbb{R}_t); P_{FA}) \leq P_D(d^{\text{GLR}}(\mathbb{R}_t); P_{FA}). \quad (3)$$

Here, $\hat{d}(\mathbb{R}_t)$ is any binary decision rule for \mathcal{H}_0 and \mathcal{H}_1 based on the residuals in (1), and $d^{\text{GLR}}(\mathbb{R}_t)$ is the *generalized likelihood ratio* (GLR) test, which asymptotically maximizes P_D . The performance of *ad hoc* methods using Monte Carlo simulations gives the value

of the left-hand side of (3), which can objectively be compared to the right-hand side.

It is of practical interest to investigate how much more efficient the GLR test is for the non-Gaussian linear model compared to an equivalent Gaussian linear model in (1), where \mathbb{V}_t has the same mean and covariance. The Gaussian distribution is a worst case distribution in that asymptotically

$$\begin{aligned} \min_d \max_{p_{\mathbb{V}_t}} P_D(d(\mathbb{R}_t); P_{\text{FA}}, p_{\mathbb{V}_t}) \\ = P_D(d^{\text{GLR}}(\mathbb{R}_t); P_{\text{FA}}, \mathbb{V}_t \in \mathcal{N}) \\ < P_D(d(\mathbb{R}_t); P_{\text{FA}}, p_{\mathbb{V}_t}^o), \end{aligned} \quad (4)$$

where $p_{\mathbb{V}_t}^o$ denotes the true non-Gaussian distribution of the noise. The gap in the inequality indicates how important it is to take the non-Gaussianity into consideration in the design. If the gap is negligible, fault detection may be designed under a Gaussian assumption, and one can hope for good performance. Otherwise, dedicated tests based on the GLR test statistic should be derived.

The GLR test involves explicit or implicit estimation of the parameter vector θ . One method is to compute the weighted least squares estimate of θ . This leads to the *best linear unbiased estimate* (BLUE) obtainable for linear systems. This is, however, not the *minimum variance* (MV) estimate for non-Gaussian noise. The MV estimator is in general nonlinear in data, and is asymptotically given by the *maximum likelihood* (ML) estimator. Relating to this, the *Cramér-Rao lower bound* (CRLB) offers a performance bound for parameter estimation (Lehmann, 1983; Kailath *et al.*, 2000), and hence indirectly also a bound for detection performance.

Of interest here is if utilizing nonlinear estimators pays off in terms of better detection performance and therefore motivates the need for more computational power. For autoregressive models the information content in noise (*intrinsic accuracy* (IA)) can be used to determine the potential for nonlinear methods (Sengupta and Kay, 1990; Kay and Sengupta, 1993). The same concept has also been used to describe optimal detectability for linear systems (Hendeby and Gustafsson, 2005). This paper elaborates further on this and points out the importance of choosing an appropriate fault basis, as this affects the detectability.

This paper is organized as follows: Sec. 2 introduces the models used, both for faults and residuals. In Sec. 3 accuracy is introduced as a measure of the information available in noise, and in Sec. 4 bounds are derived for the detection performance. Simulations in Sec. 5 exemplify the results. Conclusions are drawn in Sec. 6.

2. MODELS

To predict a system's behavior a good model is very important; the better the model, the easier it is to de-

tect abnormalities/faults. This paper assumes a linear state-space model for the system,

$$x_{t+1} = F_t x_t + G_t^w w_t + G_t^f f_t \quad (5a)$$

$$y_t = H_t x_t + e_t + H_t^f f_t, \quad (5b)$$

where y_t are measurements of the system, f_t a scalar fault, x_t the state, and w_t and e_t are mutually independent process noise and measurement noise, respectively. For detection several measurements are often considered at the same time, which can be described with a stacked model with L measurements,

$$\mathbb{Y}_t = \mathcal{O}_t x_{t-L+1} + \bar{H}_t^w \mathbb{W}_t + \mathbb{E}_t + \bar{H}_t^f \mathbb{F}_t, \quad (6)$$

where $\mathbb{Y}_t = (y_{t-L+1}^T, \dots, y_t^T)^T$ are stacked measurements, and \mathbb{W}_t and \mathbb{E}_t stacked process noise and measurement noise, respectively. The system matrices (with time dependencies removed for notational clarity) are the extended observability matrix

$$\mathcal{O} = (H^T \ (HF)^T \ \dots \ (HF^{L-1})^T)^T$$

and the *Toeplitz* matrices \bar{H}_t^f and \bar{H}_t^w

$$\bar{H}^* = \begin{pmatrix} H^* & 0 & \dots & 0 \\ HG^* & H^* & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ HF^{L-2}G^* & HF^{L-3}G^* & \dots & H^* \end{pmatrix},$$

see (Hendeby, 2005) for time-dependent expressions.

This section first shows how to structure the fault, and then construct residuals suitable for detection purposes.

2.1 Fault Model

Faults often manifest themselves in the measurements in the same way as process noise does. This introduces a difficulty because the effects of noise and faults are indistinguishable at any given time. The difference lies in the temporal behavior. Compared to noise, faults have additional structure in how they affect the system. By imposing this structure to the estimated faults it is possible to separate between noise and fault.

The fault, f_t , is split into two factors; direction and magnitude. In this paper the fault direction is contained in the matrices G_t^f and H_t^f . The remaining magnitude is then expressed as a linear regression,

$$f_t = \varphi_t^T \theta. \quad (7)$$

Written in stacked form the fault becomes

$$\mathbb{F}_t = \begin{pmatrix} \varphi_{t-L+1}^T \theta \\ \vdots \\ \varphi_t^T \theta \end{pmatrix} = \Phi_t^T \theta, \quad (8)$$

where $\Phi_t := (\varphi_{t-L+1} \dots \varphi_t)$. Above, θ is a *time-invariant* fault parameter of dimension n_θ that is independent of L , and Φ_t acts as a basis for the variation in fault magnitude. The basis should be chosen carefully to cover all faults to be detected, and at the same time

be kept as small as possible to improve detectability. Typically, $n_\theta \ll L$ to allow for efficient detection. (The effect of the size of the basis is discussed in Sec. 4.) Furthermore, with an orthonormal basis the energy in the fault is preserved in the fault parameter,

$$\|\mathbb{F}_t\|^2 = \|\Phi_t^T \theta\|^2 = \|\theta\|^2.$$

One such suitable choice for the basis is discrete *Chebyshev polynomials*, which describe orthogonal polynomials of increasing degree (Abramowitz and Stegun, 1965; Rivlin, 1974).

Example 1. (Incipient noise). Assume that a window of $L = 5$ samples is studied and that the first three discrete Chebyshev polynomials are used as basis, $n_\theta = 3$, for faults in the window. The resulting basis vectors and an example of an incipient fault is depicted in Fig. 1. \diamond




Fig. 1. The three first Chebyshev polynomials and an incipient fault described using them.

This paper adopts the convention that $\theta = 0$ indicates a fault free system. This simplifies the notation without limiting the generality.

2.2 Residual Model

To generate the residuals needed for detection, first define $\bar{H}_t^\theta := \bar{H}_t^f \Phi_t^T$, that is the effect of the structured fault in (6). This section will now show how to obtain residuals of the form (1), i.e., $\mathbb{R}_t = \bar{H}_t^\theta \theta + \bar{H}_t^v \mathbb{V}_t$, where \mathbb{V}_t contains all noise elements over the window, and \bar{H}_t^v is a coloring of that noise. In the sequel, \bar{H}_t^θ and \bar{H}_t^v are assumed thick and with full row rank, and $\text{cov}(\mathbb{V}_t) \succ 0$. For more details see (Hendeby and Gustafsson, 2005; Hendeby, 2005).

State Space with Initial State Knowledge Given an estimate of the initial conditions of a linear state-space model, \hat{x}_{t-L+1} , the residuals are given by

$$\begin{aligned} \mathbb{R}_t &= \mathbb{Y}_t - \mathcal{O}_{t-L+1}^T \hat{x}_t \\ &= \underbrace{\bar{H}_t^f \Phi_t^T}_{\bar{H}_t^\theta} \theta + \underbrace{(\mathcal{O}_t H_t^w I)}_{\bar{H}_t^v} \underbrace{\begin{pmatrix} \tilde{x}_{t-L+1} \\ \mathbb{W}_t \\ \mathbb{E}_t \end{pmatrix}}_{\mathbb{V}_t}, \end{aligned} \quad (9)$$

where $\tilde{x}_{t-L+1} := x_{t-L+1} - \hat{x}_{t-L+1}$ is the error in the estimate. The method used to estimate the state determines which distribution the error has.

Parity-Space State-Space Formulation It is sometimes unfavorable to use information about the initial state of a system. For instance, it may be impossible to get reliable information about it. An alternative is then to remove all influence on the measurements from the initial state. This is referred to as working in *parity space* (Basseville and Nikiforov, 1993). For (6) this can be achieved by multiplying the measurements with a suitable matrix, yielding the residuals

$$\begin{aligned} \mathbb{R}_t &= \mathcal{P}_{\mathcal{O}}^\perp \mathbb{Y}_t \\ &= \underbrace{\mathcal{P}_{\mathcal{O}}^\perp (H_t^w I)}_{\bar{H}_t^v} \underbrace{\begin{pmatrix} \mathbb{W}_t \\ \mathbb{E}_t \end{pmatrix}}_{\mathbb{V}_t} + \underbrace{\mathcal{P}_{\mathcal{O}}^\perp \bar{H}_t^f \Phi_t^T}_{\bar{H}_t^\theta} \theta, \end{aligned} \quad (10)$$

where by construction $\mathcal{P}_{\mathcal{O}}^\perp \mathcal{O} = 0$ and $\text{cov}(\mathbb{R}_t) \succ 0$, i.e., $\mathcal{P}_{\mathcal{O}}^\perp (H_w I)$ has full row rank. Note that these new residuals are completely independent of the initial state and the noise associated with it.

3. INFORMATION AND ACCURACY

This section introduces *Fisher information* (FI), *intrinsic accuracy* (IA), and *relative accuracy* (RA) which are important in the derivation of the fundamental performance bounds discussed in Section 4.

3.1 Fisher Information

The *Fisher information* (FI) is a measure of how much information is available about a parameter in a distribution given samples from the distribution.

Definition 1. *Fisher information* (FI) is defined (Kay, 1993), under mild regularity conditions on the distribution of ξ , for the parameter θ , as

$$\begin{aligned} \mathcal{I}_\xi(\theta) &:= -\mathbb{E}_\xi \left(\Delta_\theta^\theta \log p(\xi|\theta) \right) \\ &= \mathbb{E}_\xi \left(\left(\nabla_\theta \log p(\xi|\theta) \right) \left(\nabla_\theta \log p(\xi|\theta) \right)^T \right) \end{aligned} \quad (11)$$

evaluated for the true parameter $\theta = \theta_0$, with ∇ and Δ defined to be the *Jacobian* and the *Hessian*, respectively.

The FI is related to any unbiased estimate $\hat{\theta}$ of θ based on measurements of ξ through

$$\text{cov}(\hat{\theta}(\xi)) \succeq \mathcal{I}_\xi^{-1}(\theta) = P_\theta^{\text{CRLB}},$$

where P_θ^{CRLB} is the well known CRLB for the covariance of the estimate $\hat{\theta}$ (Kay, 1993; Kailath *et al.*, 2000) and $A \succeq B$ denotes that $A - B$ is a positive semidefinite matrix.

3.2 Accuracy

When nothing else is explicitly stated, the information is taken with respect to the mean, μ assumed to be zero, of the distribution in question, and therefore the notation $\mathcal{I}_e := \mathcal{I}_e(\mu)$, with e being a stochastic variable, will be used. This quantity is in (Kay and Sengupta, 1987; Kay, 1998; Cox and Hinkley, 1974) referred to as the *intrinsic accuracy* (IA) of the *probability density function* (PDF) for e . It follows that

$$\mathcal{I}_e = -\mathbb{E}_e \left(\Delta_\mu^\mu \log p_e(e - \mu) \right) = -\mathbb{E}_e \left(\Delta_e^e \log p_e(e) \right).$$

Theorem 1. For the intrinsic accuracy and covariance of the stochastic variable e the semidefinite inequality

$$\text{cov}(e) \succeq \mathcal{I}_e^{-1},$$

holds with equality if and only if e is Gaussian.

Proof: See (Sengupta and Kay, 1989). \square

In this variance sense the Gaussian distribution is a worst case distribution. Of all distributions with the same covariance the Gaussian is the one with the least information about its mean. All other distributions have larger IA.

To be able to easily talk about the increase in accuracy, *relative accuracy* (RA) is introduced in the following way:

Definition 2. Denote with *relative accuracy* (RA), the positive scalar Ψ_e such that $\text{cov}(e) = \Psi_e \mathcal{I}_e^{-1}$, when such a scalar exists.

It follows from Theorem 1 that, when RA is defined, $\Psi_e \geq 1$, with equality if and only if e is Gaussian. The RA is hence a relative measure of how much useful information there is in a distribution, compared to a Gaussian distribution with the same covariance.

Example 2. (Outlier distribution). Outliers in measurements can be described using a Gaussian sum,

$$e \sim (1 - \omega)\mathcal{N}(0, \Sigma) + \omega\mathcal{N}(0, k\Sigma), \quad (12)$$

where $0 < \omega < 1$ denotes how likely outliers are, k tells how much worse variance the outliers have, and Σ is the nominal measurement variance. The RA of e in (12) varies with ω and k , as depicted in Fig. 2. Note, given the right conditions a good outlier description may be much more informative than a Gaussian second order equivalent with the same mean and variance. \diamond

4. GENERALIZED LIKELIHOOD RATIO BASED DETECTION

One method commonly used for detection purposes is the *generalized likelihood ratio* (GLR) test. Given




Fig. 2. Inverse relative accuracy, Ψ_e^{-1} for (12). (\times denotes the noise used in Sec. 5.)

the PDF of the residuals conditioned on θ , $p(\mathbb{R}_t|\theta)$, the GLR test statistic is (t has been dropped for notational clarity)

$$L_G(\mathbb{R}) = \frac{\sup_{\theta|\mathcal{H}_1} p(\mathbb{R}|\theta)}{\sup_{\theta|\mathcal{H}_0} p(\mathbb{R}|\theta)} = \frac{p(\mathbb{R}|\hat{\theta}_1)}{p(\mathbb{R}|\hat{\theta}_0)}, \quad (13)$$

where $\hat{\theta}_0$ and $\hat{\theta}_1$ are the ML estimates of θ under \mathcal{H}_0 and \mathcal{H}_1 , respectively.

Using the GLR test statistic with a threshold test

$$L_G(\mathbb{R}) \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\geq}} \gamma \quad (14)$$

is then to decide if the alternative hypothesis is sufficiently more likely (defined by γ) than the null hypothesis.

This section describes the statistical properties of the GLR test and based on this conclusions are drawn about how different noises affect the detection performance.

4.1 Asymptotic GLR Test Statistics

The GLR test is known to be a *uniformly most powerful* (UMP) test among all *invariant* tests, see Theorem 2. However, there are no general results about non-asymptotic properties of the GLR test, but it is known to be optimal in special cases (Basseville and Nikiforov, 1993), and it is known that optimal tests do not exist in certain cases (Lehmann, 1986).

Theorem 2. The GLR test is asymptotically UMP, *i.e.*, most powerful for all θ under \mathcal{H}_1 , amongst all tests that are *invariant*. Furthermore, the asymptotic statistics are given by

$$L'_G(\mathbb{R}) \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\approx}} \begin{cases} \chi_{n_\theta}^2, & \text{under } \mathcal{H}_0 \\ \chi_{n_\theta}^2(\lambda), & \text{under } \mathcal{H}_1 \end{cases}, \quad (15a)$$

where $L'_G(\mathbb{R}) := 2 \log L_G(\mathbb{R})$ and

$$\lambda = (\theta_1 - \theta_0)^T \mathcal{I}(\theta = \theta_0)(\theta_1 - \theta_0). \quad (15b)$$

The dimension of $\mathcal{I}(\theta = \theta_0)$ is $n_\theta \times n_\theta$, and θ_0 and θ_1 are the true values of θ under \mathcal{H}_0 and \mathcal{H}_1 , respectively.

Proof: See (Kay, 1998, Ch. 6). \square

The important non-centrality parameter (15b) is for the class of systems studied in this paper (1)

$$\lambda = \theta_1^T H_t^{\theta T} (H_t^v \mathcal{I}_V^{-1} H_t^{vT})^{-1} H_t^\theta \theta_1, \quad (16)$$

when $\theta_0 = 0$ (Hendeby and Gustafsson, 2005).

Other tests that possess the same favorable asymptotic properties are the *Wald test* and the *Rao test* (Kay, 1998). The asymptotic performance stipulated by the theorem constitutes an upper bound for what can be achieved with a detector given finite information. According to (Kay, 1998), the performance of a GLR test is often close to asymptotic performance for relatively modest sizes of data.

4.2 Asymptotic Detection Performance

If the asymptotic GLR test described in Sec. 4.1 is used the *probability of false alarm*, P_{FA} , and the *probability of detection*, P_D , can be derived analytically. These values then constitute an upper bound on the performance obtainable for the system, due to the UMP property of the GLR test.

The P_{FA} for a given γ' is

$$P_{FA} = \Pr(\text{decide } \mathcal{H}_1 | \mathcal{H}_0) = \mathcal{Q}_{\chi_{n_\theta}^2}(\gamma'), \quad (17)$$

where \mathcal{Q}_* denotes the complementary cumulative density function of the distribution $*$. Note, P_{FA} depends only on the choice of threshold γ' and the fault parameter dimension n_θ , hence changing the noise distributions will only affect the probability of detection

$$P_D = \mathcal{Q}_{\chi_{n_\theta}^2(\lambda)}(\gamma'), \quad (18)$$

where λ is defined by (16). The function $\mathcal{Q}_{\chi_{n_\theta}^2(\lambda)}(\gamma')$ is monotonously increasing in λ (increasing the mean under the alternative hypothesis lessens the risk that a detection is missed), hence any increase in λ improves P_D . It follows directly from (16) that it is easier to detect a larger fault.

More importantly, if \mathcal{I}_V increases, P_D increases as well. Now, since any increase in RA increases λ , and since $\Psi_V > 1$ for non-Gaussian noise it follows that any non-Gaussian noise improves P_D compared to the same system with Gaussian noise. That is, if the noise is non-Gaussian this could significantly improve the ability to detect a fault if the noise is treated appropriately.

The dimension of the fault parameter is another important factor. With decreasing n_θ the variance of L'_G decreases, making it easier to tell the null and alternative hypothesis apart. Hence, the number of fault parameters should be kept as low as possible to not lose too much detection power.

5. SIMULATION STUDY

To illustrate the theory, consider a DC motor given by the state space description

$$\begin{aligned} x_{t+1} &= \begin{pmatrix} 1 & 1 - e^{-T} \\ 0 & e^{-T} \end{pmatrix} x_t \\ &+ \begin{pmatrix} T - (1 - e^{-T}) \\ 1 - e^{-T} \end{pmatrix} (w_t + f_t) \\ y_t &= (1 \ 0) x_t + e_t, \end{aligned} \quad (19a) \quad (19b)$$

with w_t and e_t mutually independent noise. Assume that the fault f_t is affine in its behavior. For the simulations in this paper: $T = 0.4$, $w_t \sim \mathcal{N}(0, (\frac{\pi}{180})^2)$, and $e_t \sim 0.9\mathcal{N}(0, (\frac{\pi}{180})^2) + 0.1\mathcal{N}(0, (\frac{10\pi}{180})^2)$, i.e., e_t has 10% outliers with 100 times larger variance. The measurement noise is then characterized by $\text{var}(e_t) = 0.0083$, $\mathcal{I}_e = 11000$, and $\Psi_e = 9.0$. Detection is performed in parity space, (10), on a window of $L = 6$ samples, a Chebyshev base is used to capture faults with affine magnitude profile, $n_\theta = 2$, and the following fault is simulated:

$$f_t = \begin{cases} 0, & t \leq 20T \\ (t - 20T)/(100T), & 20T < t \leq 30T \\ 1/10, & 30T < t. \end{cases} \quad (20)$$

With this setup the performance gain to expect by utilizing all available data for detection instead of assuming Gaussian noise is indicated in Fig. 3. The plot shows a clear potential gain from utilizing the correct noise distribution. Note that for this setup the impact of non-Gaussian process noise is only minor.

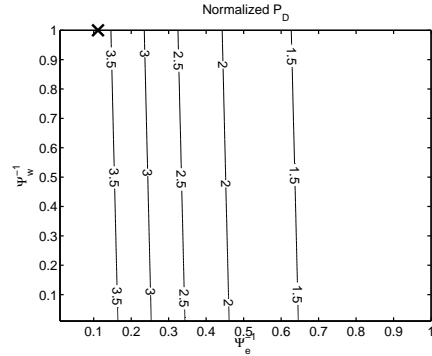


Fig. 3. Normalized P_D given $P_{FA} = 5\%$ for $f^1 = 0.1$, i.e., $\theta = (\sqrt{6}/10, 0)^T$. 1 corresponds to $P_D = 27\%$. (\times denotes the noise used in the simulations.)

Simulating the system, using a numeric ML estimate of θ in GLR tests based on the true noise distribution and also on an approximative Gaussian second order equivalent noise distribution, yields the ROC diagram in Fig. 4. P_{FA} and P_D have been derived from the first and last part of the fault sequence, i.e., $\theta = (0, 0)^T$ and $\theta = (\sqrt{6}/10, 0)$, respectively. The detection gain obtained using the true noise characteristics is impressive, especially for $P_{FA} \ll 1$. It should at the same time be pointed out that it would probably be

possible to obtain a better detector performance under the Gaussian noise approximation with modifications to the detector. However, the bound given under the approximative Gaussian assumption provides an indication of the potential performance gain.




Fig. 4. ROC plot based on 1000 MC simulations for $f_t = 0.1$, i.e., $\theta = (\sqrt{6}/10, 0)^T$.

Furthermore, the detection rate in each time instance of the simulation is presented in Fig. 5. Note how the fault is hardly detected at all with the Gaussian noise assumption whereas $P_D \approx 80\%$ when using all available information.




Fig. 5. Detection performance for a time series, when designed for $P_{FA} = 5\%$. The simulated fault included for reference. Note that different thresholds are used in the two detectors to obtain the same P_{FA} .

6. CONCLUSION

In this paper detection performance is studied for linear non-Gaussian state-space systems, for which it is shown how to construct linear residuals. It is furthermore shown how using a structured fault representation and handling non-Gaussian noise correctly may improve detection performance. Bounds for the detectability performance are provided in terms of *intrinsic accuracy* (IA) and the dimension of the fault basis used. This way optimal detection performance is related to the characteristics of the noise involved. Monte Carlo simulations, which come close to the predicted performance gain, are provided to support the theory.

ACKNOWLEDGMENT

This work is supported by VINNOVA's Center of Excellence ISIS (Information Systems for Industrial Control and Supervision) at Linköpings universitet, Linköping, Sweden.

REFERENCES

- Abramowitz, M. and Stegun, I. A. (Eds.) (1965). *Handbook of Mathematical Functions: with formulas, graphs, and mathematical tables*. Dover Publications, New York.
- Basseville, M. and I. V. Nikiforov (1993). *Detection of Abrupt Changes: Theory and Application*. Prentice-Hall, Inc.
- Cox, D. R. and D. V. Hinkley (1974). *Theoretical Statistics*. Chapman and Hall, New York.
- Hendeby, G. (2005). Fundamental Estimation and Detection Limits in Linear Non-Gaussian Systems. Lic. thesis no 1199. Dept. Electr. Eng, Linköpings universitet, Sweden.
- Hendeby, G. and F. Gustafsson (2005). Fundamental fault detection limitations in linear non-Gaussian systems. In: *Proc. 44th IEEE Conf. Decis. and Contr and European Contr. Conf.* Sevilla, Spain.
- Kailath, T., A. H. Sayed and B. Hassibi (2000). *Linear Estimation*. Prentice-Hall, Inc.
- Kay, S. M. (1993). *Fundamentals of Statistical Signal Processing: Estimation Theory*. Vol. 1. Prentice-Hall, Inc.
- Kay, S. M. (1998). *Fundamentals of Statistical Signal Processing: Detection Theory*. Vol. 2. Prentice-Hall, Inc.
- Kay, S. M. and D. Sengupta (1987). Optimal detection in colored non-Gaussian noise with unknown parameter. In: *Proc. IEEE Int. Conf. on Acoust., Speech, Signal Processing*. Vol. 12. Dallas, TX, USA. pp. 1087–1089.
- Kay, S. M. and D. Sengupta (1993). Detection in incompletely characterized colored non-Gaussian noise via parametric modeling. *IEEE Trans. Signal Processing* **41**(10), 3066–3070.
- Lehmann, E. L. (1983). *Theory of Point Estimation*. Probability and Mathematical Statistics. John Wiley & Sons, Ltd.
- Lehmann, E. L. (1986). *Testing Statistical Hypotheses*. Probability and Mathematical Statistics. 2 ed. John Wiley & Sons, Ltd.
- Rivlin, T. J. (1974). *The Chebyshev Polynomials*. John Wiley & Sons, Ltd.
- Sengupta, D. and S. M. Kay (1989). Efficient estimation for non-Gaussian autoregressive processes. *IEEE Trans. Acoust., Speech, Signal Processing* **37**(6), 785–794.
- Sengupta, D. and S. M. Kay (1990). Parameter estimation and GLRT detection in colored non-Gaussian autoregressive processes. *IEEE Trans. Acoust., Speech, Signal Processing* **10**(38), 1661–1676.