CHANGE DETECTION DESIGN FOR LOW FALSE ALARM RATES

Fredrik Gustafsson and Jan Palmqvist

Department of Electrical Engineering, Linköping University, S-581 83 Linköping

Abstract. This contribution addresses the problem of automated tuning of change detectors with given false alarm rate. By estimating a parametric distribution to the test statistics computed from real or simulated data, the threshold of the test can be computed directly. The advantage is that we can predict the threshold although there are no or very few false alarms in the used data. Using real data, the method is robust to assumed noise distributions and modeling errors. We illustrate the method on the CUSUM and GLR tests applied to friction estimation in cars and an airborne navigation system, respectively.

1. INTRODUCTION

In many change detection problems the false alarm rate is pre-specified, and the design consists in choosing the design parameters that achieve this rate and, if there is one more degree of freedom, minimize the delay for detection. The primary design parameter is often a threshold to which a test statistic is compared, and if it exceeds the threshold a change (fault) in the process is detected. Standard approaches to this design problem suffer from several problems when the false alarm rate is specified to be very low:

- In a few special cases, the test statistics have known distribution and a threshold can be taken from statistical tables. This approach is quite sensitive to modelling errors and assumed noise distributions since we are interested in the tail distribution of the test statistic, which is always sensitive to the underlying assumptions.
- We can analyze a real signal from the process under the known condition of no change, or perform Monte Carlo simulations without any changes, and then pick a threshold that gives the specified number of false alarms (namely the false alarm rate times the number of observations).Three drawbacks with this approach can be listed:
  (1) The choice of threshold is ambiguous, since we get an interval for the threshold rather than a point.
  (2) The threshold determined in this way depends only on two of the largest observations of the test statistics.
  (3) This approach is not able to predict a threshold larger than the observed test statistics for extremely low false alarm rates.

To motivate the problem of predicting thresholds outside the observation horizon, consider the navigation system in aircrafts, detailed in Section 3.

The false alarm rate should here be one in 500 hours ($100,000$ hours when GPS is used for primary means). To get a statistically well motivated threshold, we should have data, real or simulated, from at least $5,000$ hours corresponding to $18 \cdot 10^6$ samples. It is not conceivable even to simulate so many samples for this system, which has 16 states.

The approach we propose is to estimate a parameterized model of the tail distribution of the test statistics from real or simulated data. More specifically, the density function of the test statistics g_t is modeled as

$$p(g_t) = (1-\alpha)p_1(g_t) + \alpha p_2(g_t - h_0; \mu) \qquad (1)$$

where $p_2(x;\mu) = \frac{1}{\mu}e^{-x/\mu}$ is the exponential distribution and p_1 is any distribution on $[0, h_0]$. The threshold is then computed directly as

$$\hat{h} = \hat{h}_0 - \hat{\mu} \log\left(\frac{FAR}{\alpha}\right) \qquad (2)$$

where FAR is the specified False Alarm Rate. The parameter α (or h_0) is a user choice. The advantages of the method are as follows:

- It is insensitive to assumed noise distributions and modeling errors when real data are used.
- The estimated threshold is based on many observations, not only a few.
- It solves the interpolation problem for normal false alarm rates, in that we do not get an interval as an answer.
- Most importantly, we can predict thresholds outside the observation interval of the test statistics

[1] Thanks SAAB AB and AB Volvo

for very low false alarm rates.

The key problem is to choose the class of parametric distributions. The performance hinges on this choice. We will in the next two sections study two well-known representatives of change detectors, the CUSUM test and the GLR test, for which we are able to motivate the test statistic distribution and show how well it approximates a histogram over the test statistics. The result is a completely automated procedure to tune change detectors.

2. THE CUSUM TEST

2.1 *The Test*

The CUSUM test takes as it input a zero-mean white noise signal, and it gives an alarm when the mean of the signal is significantly different from zero. The CUSUM input can be any filter residual, for instance the prediction error from a Kalman filter, as illustrated in Figure 1.

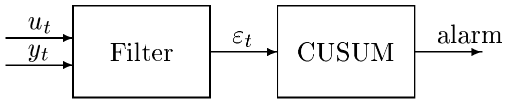

Fig. 1. The residual generator takes the observed input u_t and output y_t and transforms them to a sequence of residuals ε_t.

The CUSUM test (Page 1954, Basseville and Nikiforov 1993) is as follows:

$$\begin{aligned} g_0 &= 0 \\ g_t &= \max(0, g_{t-1} + \varepsilon_t - \nu) \\ \text{if } g_t > h & \text{ then alarm and } g_t = 0 \end{aligned} \tag{3}$$

Here the *drift parameter* ν and the threshold h are design parameters.

2.2 *Algorithm*

The distribution of the test statistic is naturally chosen as a mixture of two distributions, one being an impulse distribution at zero where we find a large number of test statistics. The other distribution is chosen as the exponential distribution. The parameter in this distribution can be initialized as the mean of the observed test statistics.

We propose the following algorithm for automatic tuning of the design parameters.

Algorithm 1. (CUSUM design). Given is the False Alarm Rate (FAR). The design parameters of the CUSUM test h and ν are chosen to achieve this FAR and minimize the Mean Time to Detection (MTD) as follows:

(1) Given a data set, real or simulated, where there are no changes. For each ν, take $h = \infty$ and apply the CUSUM test to the residuals from the first data set. Adapt the following distribution to them:

$$p_g(x) = (1-\alpha)\delta(x) + \alpha\frac{1}{\mu}e^{-x/\mu} \tag{4}$$

The threshold is then chosen as

$$\hat{h}(\nu) = -\mu \log\left(\frac{FAR}{\alpha}\right) \tag{5}$$

(2) Given a data set from the system with known changes and change times. Apply the CUSUM test and compute $\mathrm{MTD}(\hat{h}(\nu), \nu)$ and then apply a simple line search to minimize the MTD.

Note that the two data sets can be used all the time during the iterations needed in step 2, and that the filter needs to be applied only once.

2.3 *Example*

Data are here generated by the following AR(2) process:

$$y_t = \frac{1}{(1-0.8q^{-1})^2}e_t \tag{6}$$

where e_t is white noise with unit variance. Assume that the false alarm rate (FAR) is specified to 1/1000.

Simulating 100000 data, applying the CUSUM test with $\nu = 0.3$ and $h = \infty$ and applying step 1 of Algorithm 1 gives parametric and non-parametric distributions for positive g_t in the CUSUM test as depicted in the upper plot of Figure 2. As seen, the parametric distribution follows the histogram quite well.

The threshold computed from the test statistics (h_{ts}) and the one taken from the estimated parametric distribution (h_{par}) become:

$$\begin{aligned} \hat{h}_{ts} &= 10.8 \\ \hat{h}_{par} &= 11.7 \end{aligned}$$

These are close to each other and the true value because of the large number of data (compared to the false alarm rate).

The question is what happens if considerably less data are used. The lower plot of Figure 2 shows one realization using only 1000 data. That is, only one false alarm

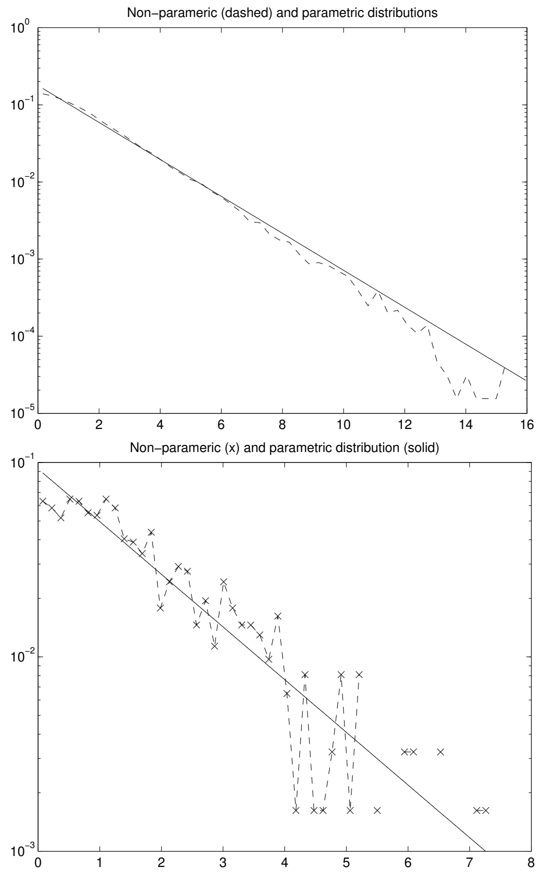

Fig. 2. Parametric and non-parametric distributions of the positive test statistics for 100000 and 1000 data, respectively.

is observed in the average. The non-parametric approach is of course quite sensitive, since it is only a function of a few of the largest observed test statistics. The parametric approach uses all test statistics larger than zero.

A Monte Carlo simulation using 50 realizations gives the following mean and standard deviation of the threshold estimates

$$\begin{aligned} \mathrm{E}(\hat{h}_{ts}) &= 8.8 \ \mathrm{Std}(\hat{h}_{ts}) &= 4.2 \\ \mathrm{E}(\hat{h}_{par}) &= 12.2 \ \mathrm{Std}(\hat{h}_{par}) &= 3.4 \end{aligned}$$

We observe that the parametric estimate is more accurate and has lower standard deviation.

2.4 *Tire–road friction monitoring*

Here an application will be described where the tire–road friction in a car is monitored for the driver. In the system, a change detector is included that alarms for sudden decreases in friction. For the drivers confidence in the system, it is of utmost importance to have a very low false alarm rate. Background and the algorithm that has been implemented in real-time is described in (Gustafsson 1997), and the change detector is detailed in (Gustafsson 1996). Basically, the algorithm consists in a Kalman filter for estimating slope and offset in a straightline model, $y(t) = \theta_1(t)\mu(t) + \theta_2 + e(t)$, and a CUSUM test applied to the Kalman filter innovations. A lot of real data are available, but here only one data sequence of the input $\mu(t)$ is used to simulate a small number of data. The simulation model is based on results from extensive test trials and verified to be accurate. The drift parameter $\nu = 0.0005$ is matched to the time constant of the Kalman filter, in that the Kalman filter takes care of small changes in friction.

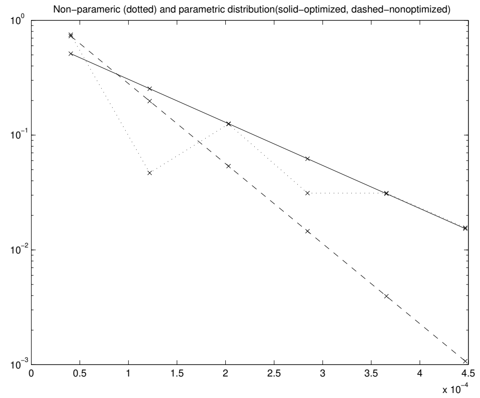

Fig. 3. Autotuning CUSUM for false alarm rate 0.01 and $\nu = 0.005$. Histogram (dotted) and parametric fit (solid, initial guess dashed).

For illustration, the false alarm rate is set to 0.01 and one single experiment is used. For a commercial system, all available real data collected on constant friction surfaces should be used, and less than one false alarm per year is acceptable. The design result of autotuning is illustrated in Figure 3. The solid line is a good model of the histogram (dotted line). The dashed line is the initial guess of parametric tail distribution. A Monte Carlo simulation over 100 realizations gave false alarm rate 0.008.

3. THE GLR TEST

With the Generalized Likelihood Ratio (GLR) Test (Willsky and Jones 1974) we have the same problem of setting the threshold to achieve specific design criterion as for the CUSUM test. In this section the GLR test for state space applications is reviewed. The expression for the test statistics is used to motivate the structure of the parametric distribution to which the GLR test statistics belong under a no change condition. Finally the proposed method is applied to a simplified real problem where the GLR test is applied for integrity monitoring of an integrated navigation system where the false alarm rate is set to be very low.

3.1 *The Test*

Consider a discrete time state space model of a system under an additive change ν in the state vector at time k:

$$\begin{aligned} \boldsymbol{x}_{t+1} &= \boldsymbol{F}_t\boldsymbol{x}_t + \boldsymbol{G}_t\boldsymbol{u}_t + \boldsymbol{w}_t + \delta(t-k)\boldsymbol{\nu} \\ \boldsymbol{y}_t &= \boldsymbol{H}_t\boldsymbol{x}_t + \boldsymbol{e}_t, \end{aligned} \tag{7}$$

Because of the linear model, the dependence of $\boldsymbol{\nu}$ on the state estimates and innovations will be linear as well. The dynamic signatures caused by this additive change can recursively be written as in (Basseville and Nikiforov 1993):

$$\boldsymbol{\varphi}^T_{x,t+1}(k) = \boldsymbol{H}_t\left(\prod_{i=k}^{t}\boldsymbol{F}_i - \boldsymbol{F}_t\boldsymbol{\mu}_{x,t}(k)\right) \tag{8}$$

$$\boldsymbol{\mu}_{x,t+1}(k) = \boldsymbol{F}_t\boldsymbol{\mu}_t(k) + \boldsymbol{K}_{t+1}\boldsymbol{\varphi}^T_{x,t+1}(k). \tag{9}$$

$$\boldsymbol{f}_t(k) = \sum_{i=1}^{t}\boldsymbol{\varphi}_i(k)\boldsymbol{S}_i^{-1}\boldsymbol{\epsilon}_i \tag{10}$$

$$\boldsymbol{R}_t(k) = \sum_{i=1}^{t}\boldsymbol{\varphi}_i(k)\boldsymbol{S}_i^{-1}\boldsymbol{\varphi}_i^T(k). \tag{11}$$

Here ε_t and K_t are the prediction errors and gain of the standard Kalman filter. From these recursions, the maximum likelihood estimate of $\boldsymbol{\nu}$, given the jump instant at k, can be written as:

$$\hat{\boldsymbol{\nu}}(k) = \boldsymbol{R}_t^{-1}(k)\boldsymbol{f}_t(k), \tag{12}$$

and the test statistics for each k is now given by:

$$l_t(k,\hat{\nu}(k)) = \boldsymbol{f}_t^T(k)\boldsymbol{R}_t^{-1}(k)\boldsymbol{f}_t(k). \tag{13}$$

The test is computed for each k giving the maximum likelihood estimation of the change time as:

$$\hat{k} = \arg\max_k l_t(k,\hat{\boldsymbol{\nu}}(k)). \tag{14}$$

A fault is declared if $l_t(\hat{k},\hat{\boldsymbol{\nu}}(\hat{k}))$ is larger than a predetermined threshold h_{GLR}.

3.2 *Algorithm*

From (13), the test statistics $l_t(k,\hat{\nu}(k))$ is a quadratic function in the residuals on the form:

$$(\boldsymbol{\epsilon}_1^t)^T\boldsymbol{A}\boldsymbol{\epsilon}_1^t,$$

where

$$\boldsymbol{\epsilon}_1^t = [\boldsymbol{\epsilon}_1^T\boldsymbol{\epsilon}_2^T\cdots\boldsymbol{\epsilon}_t^T]^T$$

with

$$\boldsymbol{\epsilon}_i \in \mathcal{N}(0,\boldsymbol{S}_t).$$

The quadratic form will in the no-fault situation be distributed (Graybill 1976) as $\chi^2(b,0)$, *i.e.* with b degrees of freedom and a non-centrality parameter equal to zero. The probability density function for a χ^2 distribution is:

$$\chi^2_x(n,0) = \left(\frac{1}{\Gamma(\frac{n}{2})2^{\frac{n}{2}}}\right)x^{(n-2)/2}e^{-n/2}$$

The second step of the GLR test is to take the maximum of all $l_t(k,\hat{\nu}(k))$. The probability distribution function of the maximum of a number of independent stochastic variables is the product of the probability distribution functions for the variables. In this case the variables are not independent, but still one can assume that the resulting probability distribution function has the same structure as the χ^2 distribution. We postulate a probability density function on the following form:

$$p_{l_t}(C,\lambda_1,\lambda_2) = Cl_t^{\lambda_1}e^{-\lambda_2 l_t} \tag{15}$$

The tail of p_{l_t} will then be distributed as in (1):

$$p_2(\mu) = \frac{1}{\mu}e^{-\mu l_t}.$$

We propose the following algorithm for automatic tuning of the threshold in the GLR test:

Algorithm 2. (GLR design). Given is the False Alarm Rate (FAR). The design parameter of the GLR test h_{GLR} is chosen to achieve this FAR:

(1) Given a data set, real or simulated, where there is no changes. Take $h_{GLR} = \infty$ and apply the GLR test to the residuals from the data set. Adapt the following distribution to the tail (starting at h_0) of the test statistics l_t:

$$p_{l_t}(x) = \frac{1}{\mu}e^{-\mu x} \tag{16}$$

(2) The threshold h_{GLR} is then chosen as

$$\hat{h}_{GLR} = h_0 - \hat{\mu}\log(FAR) \tag{17}$$

3.3 *Navigation Example*

The north horizontal channel of an integrated Inertial Navigation and Global Navigation Satellite System (INS-GNSS) is studied. The acceleration and gyro errors are assumed to consist of white noise plus a first order Gauss-Markov process for the acceleration bias and the gyro drift (Ko *et al.* 1996). The state-space description is chosen to include states for accelerometer bias, gyro drift and time correlated noise from the GNSS, giving the continuous state space equation as:

$$\begin{aligned} \dot{\boldsymbol{x}}(t) &= \boldsymbol{A}\boldsymbol{x}(t) + \boldsymbol{w}(t) \\ \boldsymbol{y}(t) &= \boldsymbol{C}\boldsymbol{x}(t) + \boldsymbol{e}(t) \end{aligned} \tag{18}$$

where

$$\boldsymbol{A} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & -\frac{g}{R} & \frac{1}{R} & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -\frac{1}{\tau_a} & 0 & 0 \\ 0 & 0 & 0 & 0 & -\frac{1}{\tau_g} & 0 \\ 0 & 0 & 0 & 0 & 0 & -\frac{1}{\tau_{gps}} \end{bmatrix}$$

$$\boldsymbol{C} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & -1 \end{bmatrix}$$

Corresponding to the following set of states:

$$\begin{bmatrix} \text{position error} \\ \text{position rate error} \\ \text{east angular error} \\ \text{acceleration bias} \\ \text{gyro bias} \\ \text{correlated measurement noise} \end{bmatrix}$$

Since high performance applications are considered small errors have been simulated for the satellite navigation system corresponding to differential GNSS or military systems. The performance data that were used in order to simulate the INS and the GNSS are given in Table 1.

	INS		
Gyro	bias	0.004	deg/hr
	corr noise	0.001	$\mathrm{deg}/\sqrt{\mathrm{hr}}$
	time const	3600	s
	white noise	0.001	$\mathrm{deg}/\sqrt{\mathrm{hr}}$
Accel.	bias	50	μg
	corr noise	5	$\mu\mathrm{g}/\sqrt{\mathrm{Hz}}$
	time const	7200	s
	white noise	5	$\mu\mathrm{g}/\sqrt{\mathrm{Hz}}$
	GNSS		
	corr noise	1.5	m
	time const	10	s
	white noise	1	m

Table 1. INS and GNSS performance used in simulations.

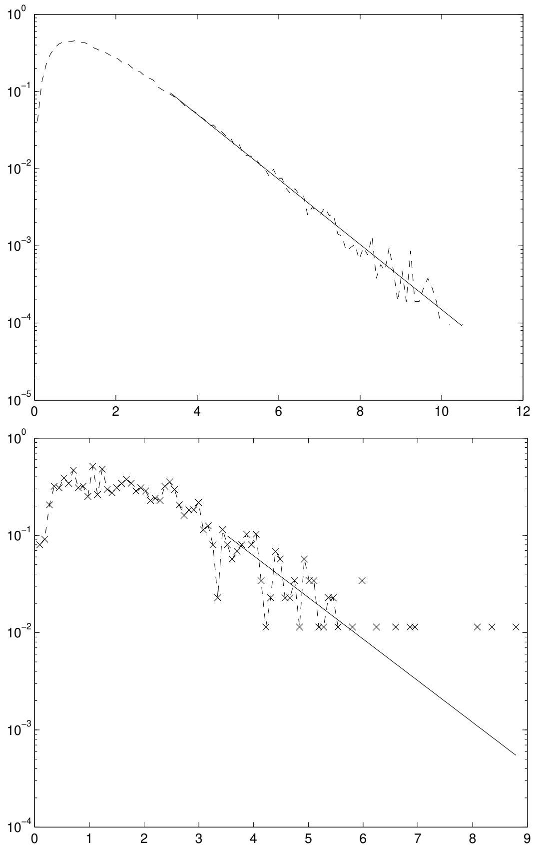

Fig. 4. Parametric and non-parametric distributions of the positive test statistics for 100,000 (upper) and 1000 (lower) data, respectively.

The Kalman filter is executed at 1 Hz. For further details on the simulation setup see (Palmqvist 1997).

In these kinds of applications the false alarm rate shall typically be set to be smaller than 0.002/hr (RTCA 1991) yielding a rate of approximately $5 \cdot 10^{-7}$ in this example.

To achieve this the Kalman filter and the GLR test (windowed over the last 10 time instants) is run for 100,000 s yielding the same number of samples. Note that this is far less than the more than 10^9 samples that should have been needed to set the threshold. To apply Algorithm 2, the upper 10% percentile of the test statistics is used, which gives an estimated parametric and non-parametric distribution for l_t as depicted in the upper plot of Figure 4. As seen the parametric distribution follows the tail of the histogram quite well.

Since we use too little data to set the threshold with the non-parametric method, a comparison is performed assuming that the FAR was set to 1/1000 and $1/10,000$ instead.

10 realizations of 10,000 data				
FAR	$E(\hat{h}_{ts})$	$Std(\hat{h}_{ts})$	$E(\hat{h}_{par})$	$Std(\hat{h}_{par})$
10^{-4}	9.21	1.13	10.62	0.63
10^{-3}	7.84	0.65	8.19	0.44
100 realizations of 1,000 data				
FAR	$E(\hat{h}_{ts})$	$Std(\hat{h}_{ts})$	$E(\hat{h}_{par})$	$Std(\hat{h}_{par})$
10^{-3}	7.08	1.09	8.08	0.91

Table 2. Results of Monte Carlo simulation of GLR threshold estimates

For FAR's of 10^{-3} and 10^{-4} respectively the thresholds computed from the test statistics (h_{ts}) and the one taken from the estimated parametric distribution (h_{par}) become:

$$\hat{h}_{ts}(FAR = 10^{-3}) = 7.87 \quad \hat{h}_{ts}(FAR = 10^{-4}) = 10.05$$
$$\hat{h}_{par}(FAR = 10^{-3}) = 8.24 \quad \hat{h}_{par}(FAR = 10^{-4}) = 10.68$$

The simulated data were divided into 10 and 100 parts of 10,000 and 1,000 data respectively and the non-parametric and the parametric thresholds were set using these data sets. In Table 2 the mean and standard deviation of the threshold estimates are summarized. One can observe that as with the CUSUM test the parametric estimate is less biased and has lower standard deviation.

If we return to the real problem and use 100,000 samples to set the threshold at $5 \cdot 10^{-7}$, Algorithm 2 would give $\hat{h}_{GLR} = 15.9$.

4. CONCLUSIONS

The problem of designing change detection tests for specified low false alarm rates was considered. An autotuning approach was presented. The key point is to adapt a parametric distribution to observed test statistics. Given the distribution of the test statistics, the threshold of the test is easily computed to give the desired false alarm rate. If there are more design parameters, an outer optimization loop was proposed for minimizing the delay for detection. We stress that only two data sets are needed, one without changes for choosing the threshold and another one with known changes for optimizing the other design parameters. This approach is more robust than approaches based on theoretical distributions of the test statistics, which rely on certain model and noise assumptions. Compared to Monte-Carlo design methods, the proposed method needs less data. The methodology was applied to the CUSUM and GLR tests. The latter was exemplified on an airborne navigation system, where the pre-specified false alarm rate is extremely low.

5. REFERENCES

Basseville, M. and I. V. Nikiforov (1993). *Detection of Abrupt Changes: Theory and Application.* Information and System Sciences Series. Prentice-Hall. Englewood Cliffs, N.J. USA.

Graybill, F. A. (1976). *Theory and Application of the Linear Model.* Wadsworth & Brooks/Cole. Pacific Grove, CA.

Gustafsson, F. (1996). Estimation and change detection of tire—road friction using the wheel slip. In: *Proc. of CCA, 1996.* Dearborn, Michigan.

Gustafsson, F. (1997). Slip-based estimation of tire - road friction. *Automatica, available at* `www.control.isy.liu.se/~fredrik/journal`.

Ko, P-Y., P. Enge and D. J. Powell (1996). Continuity Improvements via Inertial Augumentation of GPS-Based Landing System. In: *Proc. of IEEE 1996 Position, Location and Navigation Symposium.* IEEE. pp. 153–159.

Page, E.S (1954). Continuous inspection schemes. *Biometrika* **41**, 100–115.

Palmqvist, J. (1997). On Integrity Monitoring of Integrated Navigation Systems. Licentiate thesis LIU-TEK-LIC-1997:01. Department of Electrical Engineering, Linköping University. Linköping, Sweden.
available at `www.control.isy.liu.se/~janpalm`.

RTCA (1991). *Minimum Operational Performance Standards for Airborne Supplemental Navigation Equipment Using Global Positioning System (GPS).* RTCA. Washington DC, USA.

Willsky, A. S. and H. L. Jones (1974). A Generalized Likelihood Ratio Approach to State Estimation in Linear Systems Subjected to Abrupt Changes. In: *Proc. of the 1974 IEEE Conference on Decision and Control.* pp. 846–853.