A NEW SEQUENTIAL DETECTOR FOR SHORT-DURATION SIGNALS

Peter Willett

U-157, University of Connecticut
Storrs, CT 06269
willett@eng2.uconn.edu

Biao Chen

U-157, University of Connecticut
Storrs, CT 06269
bchen@eng2.uconn.edu

ABSTRACT

For quickest detection of a permanent change in distribution of otherwise *iid* observations, Page's test provides the optimal processor. Page's test has also been applied to the detection of transient (i.e. temporary) changes in distribution; it is easy to implement and has reliable performance, but as applied to the transient problem its optimality is questionable.

In this paper we offer an alternative to the Page procedure which we call the iterated generalized sequential probability ratio test, or IGSPRT. While Page's test is itself an IGSPRT, its form and performance are constrained by its reliance on constant thresholds and biases. We demonstrate that with these time-varying, markedly increased detection probabilities are possible. The IGSPRT is easiest to understand and motivate in the Gaussian shift-in-mean problem, and we discuss this in detail; but since that problem is of limited practical interest, we also examine the effect of the IGSPRT in a more realistic situation.

1. INTRODUCTION

The quickest detection of a change in distribution is an old but important problem. As an abstraction, we write

$$X_n \sim \begin{cases} dF_H(\cdot) & n < n_0 \\ dF_K(\cdot) & n_0 \leq n \end{cases} \tag{1}$$

in which $\{X_n\}$ is an independent observation sequence, n_0 is a possible and unknown change time, and where F_H and F_K are the distributions respectively before and after the change. As an example, we may consider n_0 as representing the onset of a machine's failure, and X_n the number of defective parts produced by it in the n^{th} epoch. Naturally, it is possible that no failure occurs; but if one does, it is reasonable to wish to be informed as quickly as possible, and further that the interval between false failure declarations be large.

SUPPORTED BY ONR THROUGH NUWC, NEWPORT, UNDER CONTRACT N66604-97-M-3139

The Page (or cusum) procedure [1, 2], which is optimal in this application [3] forms

$$S_n = \max\{0, S_{n-1} + g(X_n)\} \tag{2}$$

and continually tests this against a threshold h. Ideally $g(x) = \log(f_K(x)/f_H(x))$, but as long as

$$\mathcal{E}_H(g(x)) < 0 < \mathcal{E}_K(g(x)) \tag{3}$$

the scheme works. However, the value of the biasing necessary to achieve this latter is an important determinant of performance.

In the case of a transient signal, otherwise known as a temporary change in distribution, the counterpart to (1) is

$$X_n \sim \begin{cases} dF_H(\cdot) & n < n_s \text{ and } n \geq n_s + n_d \\ dF_K(\cdot) & n_s \leq n < n_s + n_d \end{cases} \tag{4}$$

Ideally one would prefer to test for such a change directly; that is, to treat this as a hypothesis-testing problem. However, since the times n_s and n_e are not known, the situation is "composite", and in general no uniformly most powerful procedure exists. As such, a generalized likelihood ratio test is often the method of choice, and it can be shown – it is not immediately obvious, perhaps – that the Page procedure provides it. An example of its application is shown in figure 1.

From this figure it would appear that the Page procedure works very well. However, in the next section we demonstrate that in the absence of information as to the transient signal's duration direct application of the Page procedure is a poor idea. We then show a new (and novel) scheme which suffers no such drawback, and should, in our opinion, be the method of choice for transient signals. Our discussion to this point will be informed by the Gaussian shift-in-mean example; in the following section we apply our scheme to a situation of greater relevance. Results are from an analytic method which is an extension of that found in [4], and are confirmed by simulation.

2. CONSTANT DETECTABILITY TRANSIENT SIGNALS

It is not unnatural that the longer the transient signal, the better the performance of Page's test. However, it must be noted that the update rule $g(\cdot)$ is in general "tuned" to a particular transient type, more specifically to a particular transient strength. While in fixed-length detection this is usually not a problem, since a procedure which works well for a low-SNR signal usually works even better for a high-SNR one, in the Page procedure this nice intuition is roiled by the fact that $g(\cdot)$ includes some implicit biasing to make certain that (3) is true.

We argue here that simply to be able to detect transients of a particular strength, and of length at least n_d, is not a useful goal. Some transient signals are short-and-loud, while some are long-and-quiet. With this in mind, we propose the following

Goal: A useful scheme ought to be able to detect any transient signal whose fixed-length detectability – that is, probability of detection in a simple hypothesis test in which start-point, end-point, and parameters are known – exceeds a given value.

We shall observe that this goal is left unsatisfied by Page's test.

By way of motivation and explanation, let us consider the very standard problem of (4) with

$$\begin{aligned} dF_H(x) &= \frac{1}{\sqrt{2\pi}} e^{-x^2/2} \\ dF_K(x) &= \frac{1}{\sqrt{2\pi}} e^{-(x-\mu)^2/2} \end{aligned} \tag{5}$$

for which the log likelihood ratio is

$$g(x) = \mu x - \mu^2/2$$

(Notice that the optimal bias $-\mu^2/2$ is already "built-in" to this update.) In this simple situation constant fixed-length detectability can be assured by specifying that

$$\mu = \frac{S}{\sqrt{\text{transient length}}}$$

meaning that the signal-to-noise ratio must remain constant regardless of the transient signal's length.

We find from figure 2 that, as indicated above, Page's test does not provide constant detectability. The Page test designed for a short-and-loud transient (denoted $L = 1, 2$ or 3) is best for such a transient; but is very poor at detecting transients of length greater than 10. Similarly, a Page scheme designed for long-and-quiet transients (denoted $L = 100$) is very poor when the length is less than 5. The reason for this anomalous behavior is revealed in figure 3, showing the design biases ($\mu^2/2$) and thresholds versus transient length. A short-and-loud transient is best served by a large negative bias and a low threshold, and is ill-served by the small bias and high threshold tuned to long-and-quiet transients; and vice-versa.

Insight is gained by considering Page's procedure as a sequence of sequential tests having thresholds zero and h. Each "reset-to-zero" of the cusum corresponds to a previous signal-absent decision, and testing begins anew. Via Wald's approximations [2] we have for a sequential test the lower and upper thresholds

$$\begin{aligned} \tau_l &= \log([1-\beta]/[1-\alpha]) \\ \tau_u &= \log(\beta/\alpha) \end{aligned} \tag{6}$$

in terms of the detection and false-alarm probabilities β and α (with $\beta > \alpha$ both small we have $\tau_l \approx 0$). This amounts, in the case of the Gaussian shift-in-mean problem, to

$$\begin{aligned} \tau_l \;\leftrightarrow\; & \textstyle\sum_{i=1}^n \log\left(dF_K(x_i)/dF_H(x_i)\right) \;\leftrightarrow\; \tau_u \\ \tau_l \;\leftrightarrow\; & \mu \textstyle\sum_{i=1}^n X_i - n\mu^2/2 \;\leftrightarrow\; \tau_u \end{aligned}$$

in which $\leftrightarrow$ denotes a threshold test. In our case the signal strength is not fixed: hence we write $\mu = \mu_n = S/\sqrt{n}$. Thus the test becomes:

$$\begin{aligned} \tau_l/\mu_n \;\leftrightarrow\; & \textstyle\sum_{i=1}^n X_i - n\mu_n/2 \;\leftrightarrow\; \tau_u/\mu_n \\ 0 \;\leftrightarrow\; & \textstyle\sum_{i=1}^n X_i - n\mu_n/2 - \tau_l/\mu_n \;\leftrightarrow\; (\tau_u - \tau_l)/\mu_n \end{aligned}$$

After substitution, we get the GSPRT (generalized in the sense of time-varying thresholds):

$$0 \;\leftrightarrow\; \textstyle\sum_{i=1}^n (X_i - b_i) \;\leftrightarrow\; h_n \tag{7}$$

for which

$$\begin{aligned} b_i &= (S/2 + \tau_l/S)\left(\sqrt{i} - \sqrt{i-1}\right) \\ h_i &= (\tau_h - \tau_l)\,\frac{\sqrt{i}}{S} \end{aligned} \tag{8}$$

are the time-varying biases and thresholds.

Motivated by this, and by the interpretation of the Page test as a sequence of SPRTs, we thus propose the iterated generalized SPRT, or IGSPRT. We do not claim that it is optimal, and in fact doubt that a feasible optimal one exists. Nevertheless, from figure 4 it is clear that the IGSPRT's performance is very good. We also offer with figure 5, in which the biases and thresholds of the IGSPRT are compared to those which would be optimal were the exact transient length known – that is, comparing the IGSPRT with a clairvoyant detector. While not identical, it is clear that the same pattern is followed in each case.

3. THE POWER-LAW STATISTIC

Nuttall has recently shown [5] that a statistic

$$T_n = \sum_{k=0}^{n-1} |X_n(k)|^{2p} \tag{9}$$

in which $\{X_n(k)\}_{k=0}^{n-1}$ denotes the DFT of the n^{th} block of data, works well (and is a reasonable approximation of optimal) at detecting signals of "unknown location, structure, extent, and strength". Useful values of p appear to be in the range $1.5 \leq p \leq 3$. The basic assumption is that $|X_n(k)|^2$ is either unit exponential (if it does not contain a signal) or exponential with mean $(1 + SNR)$ (if it does).

A Page detector based on Nuttall's statistic has already been proposed, and is analyzed in [6]. Here, we apply the IGSPRT scheme. We first note from figure 6 that the same problem – mismatch – arises here as it did in the simple Gaussian shift-in-mean case. IGSPRT biases and thresholds necessary to maintain $\bar{T} = 5\times10^6$ are plotted in figure 7. The results are shown in figure 8, which should be compared to 4 for the Gaussian case.

4. SUMMARY

In this paper we have argued that the Page procedure as applied to detection of transient signals is wanting, in that it must be "tuned" either to detect short-and-loud or long-and-quiet transients. If there is mismatch, then there is a significant chance that a detection may be missed. We have proposed a scheme – the iterated generalized sequential probability ratio test, or IGSPRT – whose biases and thresholds respectively decrease and increase as each test proceeds. Its performance appears to be good.

The IGSPRT was derived in the Gaussian shift-in-mean case, where this time-varying behavior arose naturally. In more general (and more practically-interesting) situations, such as that using a power-law pre-processor, the choice is guided but ad-hoc. Nevertheless, the performance remains good.

5. REFERENCES

[1] E. Page, "Continuous Inspection Schemes," Biometrika, vol. 41, pp. 100-115, 1954.

[2] M. Basseville and I. Nikiforov, Detection of Abrupt Changes. Prentice Hall, 1993.

[3] G. Moustakides. "Optimal Stopping Times for Detecting Changes in Distributions", Annals of Statistics, Vol. 14, No. 4, 1379-1387, April, 1986.

[4] C. Han, P. Willett, and D. Abraham, "Some Methods to Evaluate the Performance of Page's Test as used to Detect Transient Signals", submitted to IEEE Transactions on Signal Processing, June 1997.

[5] A. Nuttall, "Detection Performance of Power-Law Processors for Random Signals of Unknown Location, Structure, Extent, and Strength", NUWC Technical Report 10751, 1994.

[6] D. Abraham and R. Stahl, "Rapid Detection of Signals with Unknown Frequency Content Using Page's Test", CISS Proceedings, March 1996.

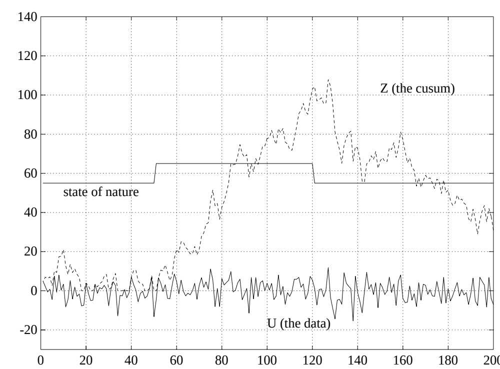

Figure 1: From samples 51 to 120 the data (U) is Gaussian with variance 25 and mean $\mu = 1$; otherwise it is Gaussian with the same variance and mean $\mu = -1$. Here Page's update uses the log-likelihood ratio.

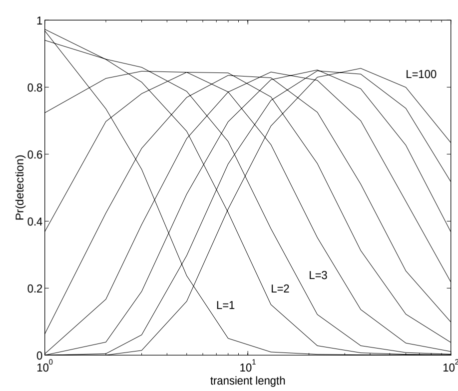

Figure 2: Detectability of Gaussian shift-in mean transients using Page procedures designed for various transient lengths L. The mean μ varies inversely with the square root of the transient's length, as described in (2). In all cases, the average time between false alarms is $\bar{T} = 5 \times 10^6$.

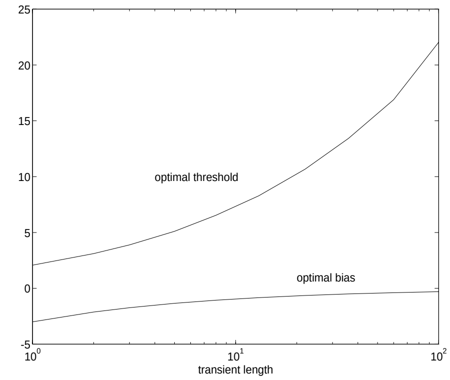

Figure 3: The optimal thresholds used in the previous figure. The bias and update for a designed transient of length L – those shown in figure 2 – can be read from this.

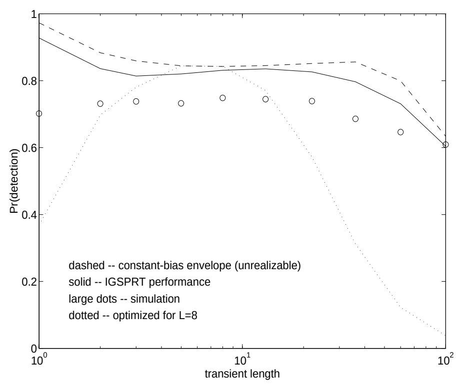

Figure 4: Performance of iterated generalized SPRT (IGSPRT) in Gaussian shift-in-mean transient problem. The performance of the Page test optimized for transient length $L = 8$ is repeated from figure 2 for comparison only

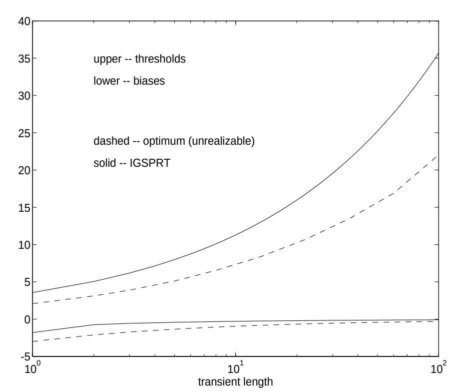

Figure 5: Comparison of thresholds used in IGSPRT to those of the Page procedure tailored to each transient length.

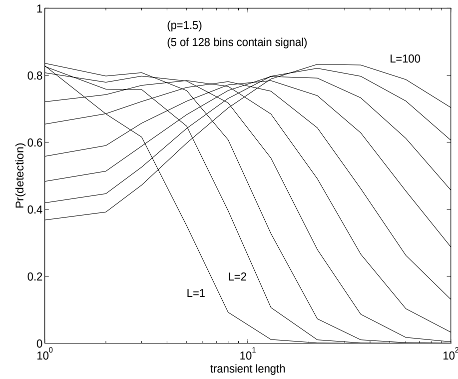

Figure 6: Power law statistic as applied with various bias factors. Here $N = 128$ (the FFT size), $p = 1.5$, $m = 5$ (5 signal-containing bins), and $SNR = 40$. The same problem as for the Gaussian case arises here.

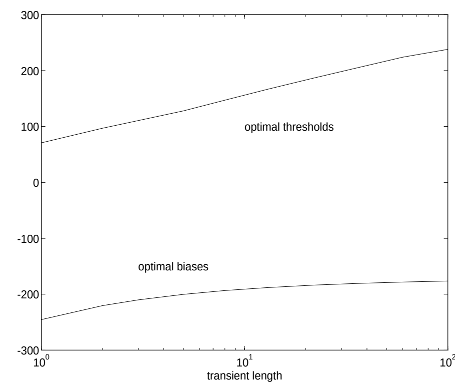

Figure 7: The thresholds and biases used by the optimal Page procedures tuned for each value of transient length for the power-law processor. Multiples of these will also be used by the IGSPRT.

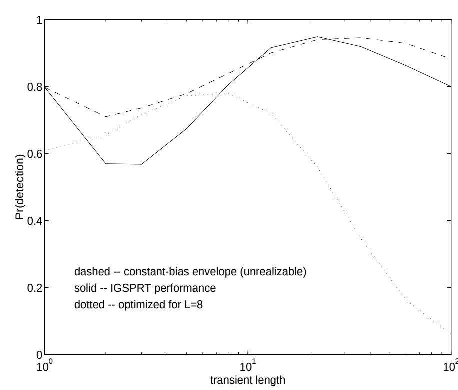

Figure 8: The performance of the IGSPRT, using as ad-hoc biases and thresholds those used by the Page tests tuned for each transient length.