SEQUENTIAL DETECTION WITH LIMITED MEMORY

Emre Ertin and Lee C. Potter

Department of Electrical Engineering
The Ohio State University
Columbus, OH 43210
ertin.1@osu.edu potter.36@osu.edu

ABSTRACT

Sequential tests outperform fixed sample size tests by requiring fewer samples on average to achieve the same level of error performance. The Sequential Probability Ratio Test (SPRT) has been suggested by Wald [1] for sequential binary hypothesis testing problems. SPRT recursively calculates the likelihood of an observed data stream and requires this likelihood to be stored in memory between samples. In this paper we study the design of sequential detection tests under memory constraints. We derive the optimal sequential test in the case where only a quantized version of the likelihood can be stored in memory. An application of the proposed techniques is large scale sensor networks where price and communication constraints dictate limited complexity devices, which store and transmit concise representations of the state of nature.

1. INTRODUCTION

In this paper we discuss the design of finite memory systems for the sequential binary hypothesis testing problem. The solution for the sequential binary hypothesis testing problem in the absence of constraints on memory is given by the SPRT. For specified error probabilities, SPRT has the property of minimum stopping time. Consequently, SPRT can be shown to be Bayes optimal minimizing the sum of observation cost and Bayes decision risk. SPRT is a stationary policy and therefore does not require a time index, but it recursively calculates the likelihood of the observed data stream and requires this likelihood to be stored in memory between observations. We consider the design of a sequential test where only a quantized version of the likelihood can be stored in memory. We model a sequential detector with finite memory as a finite state machine with state transitions directed by observations (Figure 1). Cover [2] gives an example of a statistical test where quantization of the sufficient statistics between samples leads to asymptotically suboptimal behavior and suggest that quantization for statistical tests should be approached from first principles. We derive the optimal decision rule for the binary detection problem under Bayes cost criterion and show that it is a simple partition of the a posteriori probability space specified by $L+1$ thresholds for a detector with $\log(L)$ bits of memory.

Hellman and Cover [3] consider the design of finite-memory systems for the binary hypothesis testing problem. They consider a test which reports a decision after each observation based on the current observation and what has been stored in memory. The lower bound on the asymptotic proportion of errors is derived. In general this bound cannot be attained, instead Hellman *et al.* derives a family of tests that are ϵ- optimal. Cover [2] considered the same problem for a system that can recall both the observation number and a finite memory representation of the previous measurements. Cover shows that a test with two bit statistics can achieve a limiting probability of zero under either hypothesis. Mullis and Roberts [6] also considered the design of finite-memory systems which can recall the observation number. They propose a solution for truncated sequential tests computed using dynamic programming. Solutions to two-armed bandit problems with finite memory constraints are discussed in [4, 5].

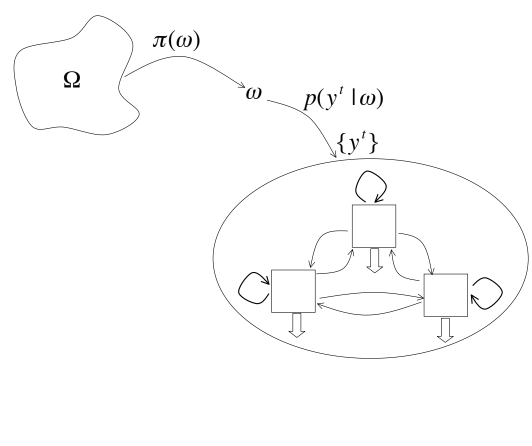

Fig. 1. Sequential detection with limited memory

2. MODEL

Let $y^0, y^1, \ldots$ be an infinite sequence of random variables in $\mathbf{R}^N$, conditionally independent and identically distributed with density $f(y|\omega)$, where $\omega \in \Omega = \{\omega_0, \omega_1\}$ represents the state of nature. Further let π be the prior probability of hypothesis ω_1 and the likelihood ratio $\Lambda(y^t)$ is defined as $f(y^t|\omega_1)/f(y^t|\omega_0)$. We assume the density functions $f(y|\omega_i)$ and the prior π are known. The sequential test is performed by a finite state machine with L states. At each state $l \in \{1, \ldots, L\}$ the decision rule is specified by the triple $(\delta^l, \gamma^l, \eta^l)$, where $\delta^l : \mathbf{R}^N \rightarrow \{0,1\}$, $\gamma^l : \mathbf{R}^N \rightarrow \{0,1\}$, $\eta^l : \mathbf{R^N} \rightarrow \{1, 2, \ldots, n\}$ specify stopping, final decision and state selection rules, respectively. If the detector is at state l it computes $\delta^l(y^t), \gamma^l(y^t), \eta^l(y^t)$ using the most recent observation. If $\delta^l(y^t) = 0$ then the detector stops and declare that hypothesis $\gamma^l(y^t)$ is true on the other hand if $\delta^l(y^t) = 1$ the detector jumps to state $\eta^l(y_t)$ and make a new measurement y^{t+1}. In this paper, we restrict our attention to nonrandomized decision rules and consider the problem of minimizing probability of error under constraints on average number of observations and memory.

3. A SEQUENTIAL TEST WITHOUT MEMORY

First consider the design of a sequential test with no memory (*i.e.*, with one state). In this case the state selection rule is trivial and we need to specify only the stopping and final decision rules $\{\delta, \gamma\}$. Define the regions $\mathcal{H}_0 = \{y \in \mathbf{R}^N | \delta(y^t) = 0 \text{ and } \gamma(y^t) = 0\}$, $\mathcal{H}_1 = \{y \in \mathbf{R}^N | \delta(y^t) = 0 \text{ and } \gamma(y^t) = 1\}$, and $\mathcal{R} = \{y \in \mathbf{R}^N | \delta(y^t) = 1\}$. If y^t is in $\mathcal{H}_i$ the detector stops and make a decision of ω_i. If $y^t \in \mathcal{R}$ the detector makes another observation. For a given test (δ, γ), we can calculate the probability of miss P_M, probability of false alarm P_{FA}, and expected number of samples $N(\omega_i)$ under hypothesis ω_i using:

$$\begin{aligned}
P_M &= \frac{P(y \in \mathcal{H}_0|\omega_1)}{1 - P(y \in \mathcal{R}|\omega_1)}, \\
P_{FA} &= \frac{P(y \in \mathcal{H}_1|\omega_0)}{1 - P(y \in \mathcal{R}|\omega_0)}, \\
N(\omega_0) &= \frac{1}{1 - P(y \in \mathcal{R}|\omega_0)}, \\
N(\omega_1) &= \frac{1}{1 - P(y \in \mathcal{R}|\omega_1)}
\end{aligned}$$

We seek to find the optimal test that minimizes the sum of expected observation cost and probability of error given by:

$$\mathbf{C} = \pi P_M + (1-\pi) P_{FA} + c((1-\pi)N(\omega_0) + \pi N(\omega_1))$$

where c is the cost of each observation. This criterion can be extended trivially to include cross terms in Bayes decision cost and state dependent cost of experimentation.

The posterior probability $\pi(y^t)$ of ω_1 after observing y^t can be computed using:

$$\pi(y^t) = \frac{\tilde{\pi}\Lambda(y^t)}{\tilde{\pi}\Lambda(y^t) + 1 - \tilde{\pi}} \;, \tag{1}$$

where $\tilde{\pi}$ is the probability of ω_1 given the fact that the test is still continuing. $\tilde{\pi}$ can be computed for a given test using Bayes rule as

$$\tilde{\pi} = \frac{\pi N(\omega_1)}{\pi N(\omega_1) + (1-\pi) N(\omega_0)} \;. \tag{2}$$

Note that the test has no memory and therefore cannot recall anything about the previous samples, not even the number of samples taken before the current sample.

Now given the test took another observation, we define the expected cost-to-go function as

$$\begin{aligned}
V(\pi(y^t)) = (1 - \pi(y^t))(cN(\omega_0)P_{FA}) \\
+ \pi(y^t)(cN(\omega_1) + P_M)
\end{aligned}$$

and the expected cost of terminating test is given by:

$$U(\pi(y^t)) = \min\{\pi(y^t), 1 - \pi(y^t)\} \;.$$

The optimal decision rules satisfy:

$$\delta(y^t) = \begin{cases} 0 & \text{if } V(\pi(y^t)) > U(\pi(y^t)) \\ 1 & \text{otherwise} \end{cases} \tag{3}$$

$$\gamma(y^t) = \begin{cases} 0 & \pi(y^t) < 0.5 \\ 1 & \pi(y^t) \geq 0.5 \end{cases} \tag{4}$$

$V(\pi)$ is linear and can intersect at most at two points with $U(\pi)$. Since the equations (1) and (2) are monotone transformations of the likelihood ratio $\Lambda(y^t)$, we have the following result.

Theorem 1 *The optimal sequential test (δ^*, γ^*) with no memory is specified by two thresholds t_0, t_1. The test is stopped and ω_0 is declared if $\Lambda(y^t) < t_0$; the test is stopped and ω_1 is declared if $\Lambda(y^t) > t_1$; and the test continues with the next sample otherwise.*

The relations in (3) and (4) are not equilibrium conditions that characterize the optimal decision rules. They are not explicit expressions for the decision rules, because the variables C_0, C_1, P_M, P_{FA} are functions of the thresholds specified in (δ, γ). A policy iteration method [7] can be used for the solution of these coupled nonlinear equations.

Example We apply the binary sequential hypothesis test without memory to constant signal detection problem in Gaussian noise for $\pi = 0.25$ and $c = 0.01$ and $f(y|\omega) \sim N(\mu(\omega), \sigma = 1)$. The resulting posterior probability of ω_1 given the test is continuing is $\tilde{\pi}_1 = 0.33$. The equilibrium cost-to-go functions and the resulting thresholds are given in Figure 2.

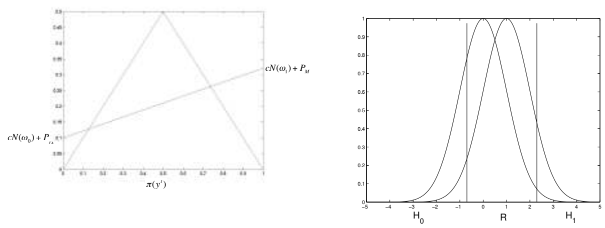

Fig. 2. Sequential signal detection without memory

4. A SEQUENTIAL TEST WITH L STATE MEMORY

Now consider the design of a sequential test using L states (*i.e.*, $\log(L)$ bits of memory). Given a test $\{\delta^l, \gamma^l, \eta^l\}_{l=1}^{L}$ we consider the resulting $L+2$ state Markov chain, where the L states are augmented by two absorbing states corresponding to stopping the test and declaring ω_0 or ω_1. We define the regions:

$$\begin{aligned} \mathcal{H}_0^l &= \{y \in \mathbf{R}^N | \delta^l(y^t) = 0 \text{ and } \gamma^l(y^t) = 0\} , \\ \mathcal{H}_1^l &= \{y \in \mathbf{R}^N | \delta^l(y^t) = 0 \text{ and } \gamma^l(y^t) = 1\} , \\ \mathcal{R}_m^l &= \{y \in \mathbf{R}^N | \delta^l(y^t) = 1 \text{ and } \eta^l(y^t) = m\} \end{aligned}$$

The $L+2$ by $L+2$ transition matrix $T(\omega)$ of the resulting Markov chain depends on the state of nature and is given by:

$$\begin{bmatrix} P(y \in \mathcal{R}_1^1|\omega) & \cdots & P(y \in \mathcal{H}_0^1|\omega) & P(y \in \mathcal{H}_0^1|\omega) \\ P(y \in \mathcal{R}_1^2|\omega) & \cdots & P(y \in \mathcal{H}_0^2|\omega) & P(y \in \mathcal{H}_0^2|\omega) \\ \vdots & & & \\ P(y \in \mathcal{R}_1^L|\omega) & \cdots & P(y \in \mathcal{H}_0^L|\omega) & P(y \in \mathcal{H}_0^L|\omega) \\ 0 & \cdots & 1 & 0 \\ 0 & \cdots & 0 & 1 \end{bmatrix}$$

The n-step transition probabilities are simply the elements of the n-fold matrix product $T^n(\omega)$. As n approaches infinity the Markov chain settles in one of the absorbing states. Therefore the limit transition matrix T^∞ has the following form:

$$T^\infty(\omega_0) = \begin{bmatrix} 0 & \cdots & 0 & (1-P_{FA}^1) & P_{FA}^1 \\ 0 & \cdots & 0 & (1-P_{FA}^2) & P_{FA}^2 \\ \vdots & & & & \\ 0 & \cdots & 0 & (1-P_{FA}^L) & P_{FA}^L \\ 0 & \cdots & 0 & 1 & 0 \\ 0 & \cdots & 0 & 0 & 1 \end{bmatrix}$$

$$T^\infty(\omega_1) = \begin{bmatrix} 0 & \cdots & 0 & P_M^1 & (1-P_M^1) \\ 0 & \cdots & 0 & P_M^2 & (1-P_M^2) \\ \vdots & & & & \\ 0 & \cdots & 0 & P_M^L & (1-P_M^L) \\ 0 & \cdots & 0 & 1 & 0 \\ 0 & \cdots & 0 & 0 & 1 \end{bmatrix}$$

where P_M^k and P_{FA}^k are respectively the probability of miss and probability of false alarm if the test is started at state k. The expected number visits to state j if the Markov chain is started at state i is given by the elements of the matrix $S_{ij}(\omega)$, which can be computed using [8]:

$$S(\omega) = \sum_{n=0}^{\infty} T^n(\omega)$$

For the particular Markov Chain corresponding to the sequential test $S(\omega)$ have the following form:

$$S(\omega) = \begin{bmatrix} N_1^1(\omega) & N_1^2(\omega) & \cdots & N_1^L(\omega) & \infty & \infty \\ N_2^1(\omega) & N_2^2(\omega) & \cdots & N_2^L(\omega) & \infty & \infty \\ \vdots & & & & & \\ N_L^1(\omega) & N_L^2(\omega) & \cdots & N_L^L(\omega) & \infty & \infty \\ 0 & 0 & \cdots & 0 & \infty & 0 \\ 0 & 0 & \cdots & 0 & 0 & \infty \end{bmatrix}$$

where $N_l^k(\omega)$ denotes the number of visits to state l for a test starting at state k under the hypothesis k. The partial row sums of $S(\omega)$ gives $N^k(\omega)$, the expected number of samples before the test terminates in one of the absorbing states if started in state k:

$$N_k(\omega) = \sum_{l=1}^{L} N_k^l(\omega)$$

In summary, T^∞ and S give the probability of detection P_M^k, probability of false alarm P_{FA}^k, and expected number of samples $N_k(\omega_i)$ under each hypothesis ω_i, given that the sequential test is started at state k. We seek to find the optimal test that minimizes the sum of expected observation cost and probability of error given by:

$$\begin{aligned} \mathcal{C} = \min_k \{ &\pi P_M^k + (1-\pi)P_{FA}^k \\ &+ c((1-\pi)N_k(\omega_0) + \pi N_k(\omega_1))\} \end{aligned}$$

Let $\pi^k(y^t)$ be the posterior probability of ω_1 given the test is at state k and y^t was observed. The posterior probability $\pi^k(y^t)$ is given as:

$$\pi^l(y^t) = \frac{\tilde{\pi}^l \Lambda(y^t)}{\tilde{\pi}^l \Lambda(y^t) + 1 - \tilde{\pi}^l} , \tag{5}$$

where $\tilde{\pi}^l$ is the probability of ω_1 given the fact that the test is at state l. $\tilde{\pi}^l$ can be computed for a given test using Bayes rule as

$$\tilde{\pi}^l = \frac{\pi N_k^l(\omega_1)}{\pi N_k^l(\omega_1) + (1-\pi) N_k^l(\omega_0)} . \tag{6}$$

Again we note that the test can only discern that it is at state l and cannot recall how it arrived to that state. Consider the decision at state l. Expected cost of continuing the test with state m is given as:

$$V^m(\pi^l(y^t)) = \pi^l(y^t)(cN_m(\omega_1) + P_M^m) \\ +(1 - \pi^l(y^t))(cN_m(\omega_0) + P_{FA}^m)$$

and the expected cost of terminating test is given by:

$$U(\pi^l(y^t)) = \min\{\pi^l(y^t), 1 - \pi^l(y^t)\} \ .$$

The optimal decision is then:

$$\begin{aligned} \delta^l(y^t) &= \begin{cases} 0 & \text{if } \min_m\{V^m(\pi^l(y^t))\} > U(\pi^l(y^t)) \\ 1 & \text{otherwise} \end{cases} \\ \gamma^l(y^t) &= \begin{cases} 0 & \pi^l(y^t) < 0.5 \\ 1 & \pi^l(y^t) \geq 0.5 \end{cases} \\ \eta^l(y^t) &= \arg\min_m\{V^m(\pi^l(y^t))\} \end{aligned}$$

Each $V^m(\pi)$ is a linear function of π. The lower envelope of the linear functions $\{V^m(\cdot)\}$ is a concave function defined with $L - 1$ intersection points. The lower envelope $\min_m\{V^m(\pi)\}$ can intersect with $U(\pi)$ at most two points. Since the equations (5) and (6) are monotone transformations of the likelihood ratio $\Lambda(y^t)$, we have the following result.

Theorem 2 *The optimal sequential test $(\delta^*, \gamma^*, \eta^*)$ with L state memory is specified by $L+1$ thresholds $t_0^l, t_1^l, \ldots, t_L^l$ for each state l. The test is stopped and ω_0 is declared if $\Lambda(y^t) < t_0^l$; the test is stopped and ω_1 is declared if $\Lambda(y^t) \geq t_L^l$; and the test continues with the next sample at state m if $t_{m-1}^l \leq \Lambda(y^t) < t_m^l$.*

Example We can interpret the optimal test as a quantization of the posteriori probability space, with the vector $\tilde{\pi}$ giving the quantization levels. Consider the binary sequential hypothesis testing problem with single bit of memory for $\pi = 0.25$ and $c = 0.01$ and $f(y|\omega) \sim N(\mu(\omega), \sigma = 1)$. Then we have $\tilde{\pi}^1 = 0.23$ and $\tilde{\pi}^2 = 0.64$. The equilibrium cost-to-go functions and the resulting decision functions for each state are given in Figure 3.

5. CONCLUSION

In this paper we derived the Bayes optimal sequential test for the binary detection problem under memory constraints. The test can be seen as an optimal quantization of the likelihood ratio. The results can be extended to sequential multiple hypotheses testing problems. The asymptotic performance of the proposed tests for large L and comparison with the SPRT [1] and M-SPRT [9] for moderate values of L are future research topics.

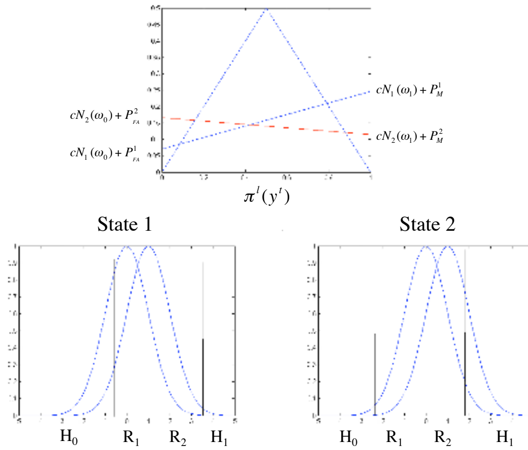

Fig. 3. Signal detection with one bit memory

6. REFERENCES

[1] A. Wald, *Sequential Analysis*, Wiley, New York, 1947.

[2] T. M. Cover, "Hypothesis testing with finite statistics," *The Annals of Mathematical Statistics*, vol. 40, no. 3, pp. 828–835, Jun 1969.

[3] M. E. Hellman and T. M. Cover, "Learning with finite memory," *The Annals of Mathematical Statistics*, vol. 41, no. 3, pp. 765–782, Jun 1970.

[4] T. Cover and M. Hellman, "The two-armed bandit problem with time invariant finite memory," *IEEE Trans. on Information Theory*, vol. 16, 185-195 1970.

[5] T. Cover, "The two-armed bandit problem with finite memory," *Information and Control*, vol. 12, pp. 371–377, 1968.

[6] C. T. Mullis and R. A. Roberts, "Memory limitation and multistage decision process," *IEEE Trans. on Systems Science and Cybernetics*, vol. SSC-4, no. 3, pp. 307–316, Sep 1968.

[7] D. P. Bertsekas and J. N. Tsitsiklis, *Neuro-dynamic Programming*, Athena Scientific, Bellmont, MA, 1996.

[8] G. C. Pflug, *Optimization of Stochastic Models*, Kluwer Academic, 1996.

[9] C. W. Baum and V. V. Veeravalli, "A sequential procedure for multihypothesis testing," *IEEE Transactions on Information Theory*, vol. 40, no. 6, pp. 1994–2007, Nov 1994.