

**A STRATEGY FOR CLASSIFYING A SET OF
DISSIMILAR CHANNELS BY THEIR A PRIORI
CHANNEL OCCUPANCY PROBABILITY**

BY SHRIDATT SUGRIM

**A thesis submitted to the
Graduate School—New Brunswick
Rutgers, The State University of New Jersey
in partial fulfillment of the requirements
for the degree of
Master of Science
Graduate Program in Electrical and Computer Engineering**

**Written under the direction of
Predrag Spasojevic and Melike Baykal Gursoy
and approved by**

New Brunswick, New Jersey

January, 2014

ABSTRACT OF THE THESIS

A Strategy for Classifying a Set of Dissimilar Channels by Their A Priori Channel Occupancy Probability

by Shridatt Sugrim

Thesis Director: Predrag Spasojevic and Melike Baykal Gursoy

Recent Changes in policy regarding the opportunistic use of licensed radio spectrum have paved the way for new innovative technologies like cognitive radio (CR). In CR systems a secondary user (SU) is allowed to use open channels if the primary user (PU) is not currently using them. Regulatory bodies like the FCC establish maximum interference requirements for SUs when making use of these channels. To comply with these requirements SUs must measure the occupancy of each of the channels they intend to use. Any strategy employed for opportunistic spectrum usage has to consider the tradeoffs between time spent searching for empty channels and time spent using those empty channels.

In most cases the spectrum sensing that is employed by a CR system starts with no prior information about the occupancy of the channels it intends to use. We propose a novel method of addressing this lack of prior knowledge by employing an efficient strategy that classifies some of the channels the SU intends to use within a fixed time limit. This classifier can be run before the SUs attempt transmission, and will provide the SUs' spectrum sensing sub-systems with a set of occupancy probabilities for some of the channels.

Our classifier will be designed around sequential probability ratio tests (SPRT)

because these tests maintain bounds on classification errors while using the smallest number of samples for classification. The classifier will attempt to classify as many channels as possible within the given time limit by intelligently allocating channel measurements. We will examine the system's performance in various measurement resource regimes and identify regimes where our approach is superior to simpler classification schemes.

Table of Contents

Abstract	ii
List of Figures	vii
1. Motivation	1
2. Proposed Model	4
3. Related work	6
3.1. Bandit Problems	6
3.2. Use of SPRT	8
4. Sequential Testing for a Single Channel	10
4.1. Composite Hypothesis and the UMP	10
4.2. More than 2 Hypotheses	13
4.3. Definition of Sequential Test	13
4.4. Finding A and B	14
4.5. Types of tests	16
4.6. Fundamental Curves	18
4.7. Derivation of $L(p)$	19
4.8. Curve plots	20
4.9. Threshold lines and Simulation	22
4.10. Random walks between two thresholds	23
4.11. Computing $\beta_{effective}$	24
4.12. $\alpha_{effective}, \beta_{effective}$ and δ	26
4.13. Multiple SPRTs	28

4.14. $L(p)$ and $E_p(M)$ for two-tests	30
4.15. Parameter sensitivity	32
5. The Cost Classifier	35
5.1. Cost Definition and Greedy Allocation	35
5.2. Distance metrics	37
5.3. Early mistakes	38
5.4. Conditions for Optimality of the greedy approach	40
5.4.1. When is greedy approach optimal for our search	40
5.4.2. Why greedy is not always optimal	41
5.4.3. Lower bound on the cost to characterize a single channel	42
5.4.4. Choices of λ that bring the search closest to optimal	42
6. Simulation and Analysis	44
6.1. Validating our implementation of the SPRT	44
6.2. Choosing δ for the SPRT	45
6.3. Effects of ζ	47
6.4. Optimal λ for the search	48
6.4.1. λ effect on searching for channels	49
6.4.2. How often do I pass on channels I should have stayed on?	51
6.5. Alternative schemes for comparison	52
6.5.1. The simplest allocation scheme	52
6.5.2. A tree structure approximation	54
6.5.3. Parameter Choices for the tree approximation	57
6.5.4. Alternative tree approaches	58
6.6. Performance Metrics	59
6.6.1. Error Probability Across the entire search	60
6.6.2. Number of Channels classified	61
6.7. Comparison of performance between different allocation schemes	63
6.8. Measurement Distribution	65

7. Conclusion	68
7.1. Why a proof of optimality is hard	68
7.2. Possible Applications	69
7.3. Future Work	70
7.4. Final Remarks	70
References	72

List of Figures

4.1. Operator Curve	21
4.2. Expected sample length curve	22
4.3. Example of the Test	23
4.4. Random Walk between two thresholds	24
4.5. The over shoot of $\beta_{effective}$	25
4.6. δ vs β	26
4.7. Large delta (big indifference region) expected sequence length	27
4.8. Shift of the indifference region boundaries due to ζ	28
4.9. Variance of a Bernoulli Random Variable	29
4.10. Sample run of 3-ary test with 2 sets of threshold lines	30
4.11. Expected Sequence length, $E_{2,p}(M)$, for a combined two-test	31
4.12. The $L_2(p)$ OC curve	33
4.13. Operator curves with fixed p' and varying δ	33
4.14. Expected sequence length with fixed p' and varying δ	34
5.1. Graphical representation of outer threshold distance calculation	38
6.1. Single Test threshold lines for varying delta.	45
6.2. Effective error rates for varying delta.	46
6.3. $max_p(E_p(M))$ for varying delta.	46
6.4. $L(p)$ for varying ζ with $\delta = 0.125$	47
6.5. $E_{p_n}(M)$ for varying ζ with $\delta = 0.125$	47
6.6. Threshold lines for varying ζ with $\delta = 0.125$	48
6.7. Samples required to find one low $E_{p_n}(M_n)$ channel and Samples required to categorize that channel vs λ	50
6.8. Actual value of the paramter p_n found vs λ	51

6.9. Channels passed on as a function of λ and δ	53
6.10. Number of channels found for all 3 schemes as functions of N	63
6.11. β_{search} for all 3 schemes as functions of N	64
6.12. Number of channels found for greedy-SPRT vs Tree schemes as functions of N	64
6.13. β_{search} for greedy-SPRT vs Tree schemes as functions of N	65
6.14. Number of channels passed on as a function of N	66
6.15. Mean samples spent classify vs searching as functions of N	66
6.16. Mean number of channels that receive a given level of sampling	67
7.1. CPU Time for All 3 schemes as function of N	71

Chapter 1

Motivation

Radio spectrum is a limited resource that is currently under heavy contention. Regulatory bodies like the FCC take on the daunting task of fairly distributing this resource among the many spectrum hungry users. Recent changes in policy regarding the opportunistic use of licensed spectrum has paved the way for new innovative technologies like cognitive radio [1].

This adaptive radio policy allows secondary users (SU) to use spectrum not allocated to them provided that the primary user (PU) is not currently using it. A key element of this technology is sensing whether the primary user is present or not. This sensing process typically requires a very extensive sensing period because of the requirement that the interference that the PU sees be minimized. At the same time, however, if the sensing takes too long, the utility of secondary usage of the spectrum goes down rapidly. If a sensing system takes too long to converge on a transmit strategy, it may miss spectrum opportunities or fail to meet transmission quality requirements.

Since the radio is adaptive, an on-line sensing plan with small convergence times is required to optimally utilize spectrum. This area of research is very active as there are several proposed strategies for sensing the environment. Techniques such as C-SPRT, and finite horizon dynamic programming are used to suggest optimal stopping times for sensing and balance points between exploration and exploitation. Many of these strategies however assume a prior distribution that is either unknown or uninformed. All of these strategies could benefit from any amount of *a priori* information about the occupancy of the channels they intend to use.

The main goal of these sensing strategies is to come up with estimates of channel occupancy that can then be used to plan a transmission pattern. Their performance can

be measured in achieved throughput. In this work we propose a different goal meant to augment these existing systems, the goal of spectrum classification.

Spectrum classification can help reduce convergence times and facilitate efficient spectral usage. Since data usage patterns are coupled with the daily routine of users, classification based on geography and time of day will be of high utility. However practical approaches to spectrum classification are not very well explored. Classification can possibly be done by two different approaches, an infrastructure approach and a mobile approach.

In the infrastructure case measurement opportunities are infinite, but mobility is not possible. In this scenario we still need to maintain a bound on error. While we do not need to be as conservative with our measurement opportunities, we would like to have occupancy information as early as possible. Our technique prioritizes channels based on their ease of classification, and can return decisions as they become available without having to wait for the entire search to complete. In the mobile case, we are not tied to fixed geographic locations, however because we are mobile we have a limit on the amount of measurement resources per geographic location. This suggests the need for measurement strategies that maximize some utility metric under a constraint on sample size / error rate.

The mobile measurement case can arise from regulatory bodies doing characterization sweeps across geographic areas, to assess usage and possibly coordinate SUs via a control channel. It could also be done by mobile SUs, in an attempt to build a personal geographic map of the occupancy for the SU. This map can be consulted before the initial spectrum scan that precedes a transmission. Map usage can lower convergence time the initial sensing, as well as warn of recurrent historical spectral artifacts that might be coming up. If the latter is done in some form of a handset the measurement resources might be very limited because of battery conservation policies, or other resource constraints.

In the following work, we propose a novel strategy for distributing a limited number of measurements amongst a set of channels with unknown occupancy probability. Our goal is to classify as many channels as we can with the limited measurement resources

while maintaining a bound on the probability that our classifier makes a mistake. The classification will partition a set of channels into subsets with similar probability of occupancy. This classification can be used to adjust future sensing policies, perhaps by starting with channels that are in the set of low occupancy channel first.

Chapter 2

Proposed Model

For the model proposed, we consider a set of N channels, where $|N|$ denotes the cardinality of this set, i.e., the number of channels to be searched. We assume that there are L possible measurements and that the rate of measurements is fixed (1 measurement per unit time). The number of channels may be greater (possibly much greater) than the number of measurements. Each measurement can correctly identify if a channel is busy or not. We also assume that each channel n is measured at time l independently of the others. Channel n is occupied with probability p_n which comes from a uniform distribution. Hence, each measurement is considered to be a Bernoulli indicator random variable $B(p)$ with 1 indicating occupied, and 0 indicating free. A complete characterization for a block of time with L measurements would be a set of \hat{P}_N of $|N|$ probabilities, one for each channel, $n \in N$. These $\hat{p}_n \in \hat{P}_N$ are estimates of the probability of being occupied, p_n . In most cases of interest we will not be able to cover the entire set of channels in the allotted time block.

Our knowledge of occupancy can be measured by splitting N into 2 sets N_{class} and $N_{unclass}$, the set of classified and unclassified channels. If a channel n is in N_{class} , then it has been measured enough times that we expect an error probability of the classification to be bounded by maximum determined by our design parameter choices. In general, the number of measurements required to make a classification will not be uniform over the set of channels. All channels that are not fully classified at the end of a time block are deemed unclassified, that is they are in $N_{unclass}$.

The result of our strategy should be a sequence of channel measurements $(l_1, l_2, \dots, l_j, \dots, l_L)$ with $j \in [0, L]$ that will maximize $|N_{class}|$ while maintaining a bound on probability of error. Each l_j is the index of the channel to be measured (that is, one of the ns).

Determination of this sequence will be dictated by our allocation strategy which will depend the ratio of $\frac{L}{|N|}$ as well as the desired error tolerance. We will introduce the sequential probability ratio test (SPRT) as our method of testing composite hypothesis on a single channel. We will then use the sequential nature of the test to propose a cost metric that will guide allocations to channels that complete their classification in fewer samples.

Chapter 3

Related work

The broad category of spectrum sensing for cognitive radio is actually quite well explored. Several surveys like [2] have been published since Mitola coined the term in 1999 [3]. In [2], Yucek charts the various strategies for sensing spectrum ranging from simple energy detection to multi-dimensional techniques. All of these techniques employ some manner of detection theory to control their error rates for the purposes of collision avoidance and efficient spectrum utilization. The key distinction between our work and many of the methods suggested in [2] and many similar surveys ([4],[5],[6]), is that their primary goal is maximizing throughput by efficiently finding spectral holes. In almost every case their primary metric is maximum throughput under different constraints, e.g. fair access to all users. In our proposed method, on the other hand, we will explore different performance metrics for an entirely different class of problems, that of classification of channels.

In the broader literature SPRTs have been used to classify students taking adaptive tests, work pioneered by Spray in 1993 [7]. A modern treatment of the techniques was given in [8]. The use of classification techniques to build priors seems unexplored in the literature. While there is a rich base of spectrum analysis techniques to draw from and the process of making classifications by using SPRTs has been well studied, merging the two disciplines to “prime” our knowledge of the spectral landscape seems novel.

3.1 Bandit Problems

In [9] and [10], Lai and Poor proposes an adaptive allocation model for sensing that employs solutions to multi-armed bandit problems. In the single user case, he uses a network model very similar to the proposed model of our work. In the discussion, he

considers a time slotted collection of channels. At each time slot, the SU must first listen to the channel. If the channel is empty, the SU will transmit. If not the SU will update their occupancy model to avoid picking this channel again. He defines a performance metric given by

$$W_{\Gamma} = E_{f(\theta)} \left\{ \sum_{j=1}^T B Z_{s(j)}(j) \right\},$$

where W_{Γ} is the expected throughput for a channel selection strategy $\Gamma = \{s(1), s(2), \dots, s(T)\}$. $s(j)$ is choice of channel at time j , and B is channel bandwidth, a constant. $Z_{s(j)}(j)$ is an indicator random variable that equals 0 if the channel is occupied, and 1 if it is not. Finally $f(\theta)$ is the joint distribution on channel parameters $\theta = (\theta_1, \dots, \theta_N)$. For a choice $s(j) = i$, $i \in [0, \dots, N]$, θ_i is the probability that channel i is unoccupied, that is $Z_i(j) = 1$. Lai argues that if θ is known, this problem is trivial, simply choose i that maximizes $\theta_i \forall j$.

If the θ is not know, then this problem can be cast as a bandit problem where we assume some distribution on the parameters, and proceed by an inductive procedure that runs the choices backwards. At each phase the author is selecting the previous step that leads him to the current maximal value of W_{Γ} , a reverse inductive process. In the finite Horizon case there is an absolute stopping time dictated by the problem it self.

Our approach has a very different utility function that aims to limit the amount of time spent sensing any individual channel. In the bandit formulation if a channel is unoccupied through out the duration of our sensing interval, this channel is always the best choice. However in our problem, if a given channel is unoccupied it is only of limited utility. An unoccupied channel will require a small number of samples to confirm that it is indeed unoccupied. Once we have determined that it is empty with a sufficiently small probability of error, this channel is decided and thus sampling it more is of little utility to our goal.

3.2 Use of SPRT

While many techniques fall back on the classical Neyman-Pearson methods of error analysis, a few, most notably Xin and Lai in [11], employ sequential testing techniques. Here the authors focus on a very similar class of problem, an attempt to minimize a cost function that depends on the expected sequence length and the error probability. Their probability model however is very different from the model we're working with. Their model comes from noise power observations where the signal is required to be present during the entire observation period. Specifically, they are trying to decide between $Y_j^{(k)} \sim \mathcal{CN}(0, [\sigma^{(k)}]^2)$ and $Y_j^{(k)} \sim \mathcal{CN}(0, P^{(k)} + [\sigma^{(k)}]^2)$, where j is the time index, and k is the channel index. We do not assume a signal model that requires constant signal, instead we expect that an underlying detector can hand us a instantaneous binary decision as to whether the channel is occupied or not. Our channels can be partially occupied with occupancy probability p_n .

The error terms in their system are a function of the current value of the estimate of the probability of occupancy. The authors employ an expected cost to go function $\tilde{J}_{j,T}(F_j)$ where T is the stopping time (given), F_j is the set of observations at time j . The \tilde{J} function is defined recursively via a reverse induction. At each step the author selects the channel that minimizes the expected cost to go, starting from the deadline and working backwards. The author then proposes several methods on how to truncate the sequence in order to use fewer samples.

We approach the choice of channels using very different strategies. Our approach attempts to minimize cost by making greedy choices, while they try to minimize an expectation across all channels. A key difference in our method is that we do not rely on reverse computation of the \tilde{J} function. Instead we use the test itself to project the best case of expected sequence length, given the current observations. If this best case is more than an unexplored channel, we move on.

It is noted that in this scenario higher probability of false alarm, P_{FA} , results in lower spectrum utilization, and higher probability of missed detection, P_{MD} , results

in more collisions with the PU. If an SPRT is not allowed to reach a complete decision (by truncating it), the bound on error that the test prescribes will not be met. Truncating the test is equivalent to running a test with looser thresholds. We show through explicit calculation, that when the thresholds are relaxed, the effective error rate is actually higher than the targets set when designing the test, see section (4.11). In our approach we treat the error rates as upper bounds, and then deem parameter set for the SPRT admissible if they achieve an error rate below these bounds. We then select the admissible parameter set that yields the largest set of completely categorized channels.

Chapter 4

Sequential Testing for a Single Channel

We build out the statistical machinery to classify a single channel. We will evaluate the classical hypothesis testing strategy and then turn to the SPRT as a test that minimizes the required samples while maintaining a specific probability of error. We will then try to find a good choice of parameters for our case.

4.1 Composite Hypothesis and the UMP

The classical binary hypothesis testing problem tries to decide between two simple hypotheses, H_0 versus H_1 . These hypotheses pertain to a generic parameter (or set of parameters) θ which controls the distribution on the general random variable we wish to make inferences about X . Here X is the RV that models the phenomenon we wish to observe, i.e., channel occupancy. X may not be a single random variable, but a vector of fixed length $\mathbf{X} = (X_1, X_2, \dots, X_L)$, where L is a fixed limit on the number of samples we can take. θ may be one of many parameter choices, that is we can consider the family Θ of which θ_0, θ_1 are members. Our hypotheses then takes on the form:

$$\begin{aligned} H_0 &: \theta = \theta_0, \\ H_1 &: \theta = \theta_1. \end{aligned} \tag{4.1}$$

These hypotheses are deemed simple because the parameter is constrained to two values due to the equality of the defining equation (we are isolating two from the family Θ). The solution to this problem is usually found by comparing a function of the observations of the random variable to a threshold η . Typically this function is the likely hood ratio defined as:

$$\Lambda(x) = \frac{F(H_0|x)}{F(H_1|x)} \geq \eta. \tag{4.2}$$

where we declare H_0 when $\Lambda(x) > \eta$ and reject H_0 when $\Lambda(x) < \eta$. As a side effect note that if $\theta > \theta_0$, this test will decide H_0 because the observations of x that would confirm H_0 become more likely that our tests expects them to be. The exact opposite happens if $\theta < \theta_1$. The choice of the threshold η can be arrived at by many different methods, see [12]. Tests of this form are named the likelihood ratio test, or LRT for short.

A general solution to any decision problem is to form a rule $\phi(x)$, called the decision rule, that is a function of the observation, x , which is chosen to minimize some cost or error probability. The function $\phi(x)$ assigns a decision to values of observable X . Before it is observed X is a random variable. Since $\phi(X)$ is a function of a random variable, it is also a random variable. For an arbitrary hypothesis test, there are two types of errors that can be made, both of which result in declaring a specific hypothesis when the other is true. In the specific case of the threshold tests which are decided by computing $\Lambda(x)$, these errors arise because the observation x causes $\Lambda(x)$ to incorrectly cross the threshold. The probability that such a crossing occurs can be computed from the observation's distribution conditioned on the correct hypothesis, e.g., $P_X(\text{declare } H_1 | \text{given } H_0 \text{ is true}) = P_X(\Lambda(x) \leq \eta | \text{given } H_0 \text{ is true}) = \alpha$ and $P_X(\text{declare } H_0 | \text{given } H_1 \text{ is true}) = P_X(\Lambda(x) \geq \eta | \text{given } H_1 \text{ is true}) = \beta$.

Since we plan to classify channels based on their parameter, we will need to deal with the entire parameter family Θ . This will require that we use composite hypothesis instead of simple equalities, we will need to consider hypothesis of the form:

$$\begin{aligned} H_0 & : \theta \geq \theta_0, \\ H_1 & : \theta \leq \theta_1. \end{aligned} \tag{4.3}$$

where we leave open the possibility that $\theta_0 = \theta_1$.

To treat this class of problems we can appeal to the general formulation of the solution as derived in [13]. Consider an arbitrary $\phi(x)$ as the probabilistic rule for deciding H_0 as a function of the observation x . Since the observation space $|X|$ is of fixed, the probability that $\phi(x)$ will decide H_0 across the set of possible observations is

given as

$$P(\text{declare } H_0)_\theta = E_\theta(\phi(X)) = \int \phi(x) dP_\theta(x). \quad (4.4)$$

As we vary θ this expectation gives us the probability of declaring H_0 as a function of θ . In the region where H_0 is wrong, we have $\beta(\theta) = P(\text{declare } H_0)_\theta \forall \theta \leq \theta_1$, the probability of making the wrong decision given θ . This integral should decrease as θ increases away from θ_1 because the probability of observing x that confirm H_0 decreases. $P(\text{declare } H_0)_\theta$ turns out to be a very important quantity for designing tests, we name it the operator characteristic (OC) curve. The same method can be used to compute an expression for $\alpha_\phi(\theta)$.

Each choice of ϕ partitions the observation space into two sets, X_0 and X_1 . Observations $x \in X_0$ confirm hypothesis H_0 , and equivalently $x \in X_1$ confirm hypothesis H_1 . The integrals $\alpha_\phi(\theta)$ and $\beta_\phi(\theta)$ are coupled via this partitioning. If we modify ϕ by moving an observation from X_0 to X_1 , the value of $\beta_\phi(\theta)$ decreases because there is one less point that $\phi(x)$ will declare H_0 on. At the same time $\alpha_\phi(\theta)$ gains that point and thus increases slightly. The opposite occurs when we move from X_1 to X_0 .

The observation space however admits a natural partitioning given as X'_0 , X'_1 , and X'_I . Here observations that fall into X'_0 and X'_1 should properly confirm H_0 and H_1 respectively. The observations in X'_I are indifferent. It is not clear which hypothesis they confirm, however our rule $\phi(x)$ has to make a decision on them. How ϕ allocates the elements of the indifference region will ultimately dictate the values of $\alpha_\phi(\theta)$ and $\beta_\phi(\theta)$.

A typical approach that appeals to the Neyman-Pearson lemma is to fix a value of α , called the size of the test, and then choose a $\phi(x)$ that maximizes $1 - \beta$, called the power of the test. If a single choice of ϕ works across the entire Θ space, this choice is said to be the uniformly most powerful test. We can use the UMP to resolve the composite hypothesis problem as it will have bounded error for all θ .

The UMP test has two major draw backs. The first is that it may not always exist, however in many simple cases a UMP can be found and used as a flawed test for more complex distributions. The second drawback is that the assumption that the size of L

remains fixed means we cannot adjust the number of samples we take without having to choose ϕ . This presents a very big problem for us because we have no mechanism for determining what a minimum number of samples should be, or if we have reached it.

While the UMP will not serve our purposes, it does provide some insights into a possible solution. The idea that a test intended to evaluate a simple hypothesis can be used to test parameters it was not explicitly intended to test as long as we have a bound on the error, will play a crucial role in the test we will use. In [14], it is shown that the sequential probability ratio test (SPRT), under the very weak conditions, will achieve the lowest expected sequence length. This test will serve as the basis for our work.

4.2 More than 2 Hypotheses

There is no specific limitation on the number of hypothesis we can decide between. If we allow for more than two hypotheses, we can use a similar definition of $\Lambda(y)$ to decide between each of these hypotheses. Let Γ be the number of hypotheses. This $\Lambda(y)$ however will now be computed across pairs of hypothesis, there should be a Λ for each pair of tests for a total of $\Gamma - 1$. If we select a minimum error criteria, we can carve up the space of possible observations into Γ sets, where each allocation is chosen to minimize the overall error. Computation of this error, however, now becomes more difficult because each hypothesis has $\Gamma - 1$ alternatives. There are some well known approximations for this error probability, e.g. the union bound.

One question becomes how to make the actual decision, [15] outlines a tournament procedure that picks hypothesis by testing a sequence of binary hypothesis elimination rounds. In the upcoming sections we will take a similar approach to decide between multiple composite hypothesis.

4.3 Definition of Sequential Test

To begin consider a single parameter simple hypothesis where we decide between two values of the generic parameter θ for the general distribution $f(x, \theta)$. Suppose the hypothesis is:

$$H_0 : X \sim f(x, \theta_0), \quad (4.5)$$

$$H_1 : X \sim f(x, \theta_1). \quad (4.6)$$

$f(x, \theta)$ could be a discrete Bernoulli distribution with parameter p , or a continuous Gaussian distribution with known variance, and mean θ . We can specify the form of the SPRT as

$$B < \Lambda'(x) = \frac{f(x_1, \theta_1) \dots f(x_m, \theta_1)}{f(x_1, \theta_0) \dots f(x_m, \theta_0)} < A. \quad (4.7)$$

This is a departure from the classical likelihood ratio defined in section (4.1) in that we specify two thresholds and do not require a fixed observation size. The SPRT is a special case of the LRT where the LRT will also be a function of m , the sample index for this test. As m grows, if the LRT goes above A , we declare H_1 , if it goes below B , we declare H_0 , otherwise we keep sampling. In [16] it is shown that there exists some maximum M for which this test terminates in a decision, that is this LRT will eventually cross A or B , i.e., $M = \inf\{n > 0 \mid \Lambda'(x) < B \text{ or } A < \Lambda'(x)\}$. Hence, M is the random stopping time.

4.4 Finding A and B

To determine the values of A and B , we first need to establish some important inequalities. Consider a sequence of observations x_1, \dots, x_M that results in the declaration of H_1 at the M^{th} step. Then we have

$$B < \frac{p_{1M}}{p_{0M}} = \frac{f(x_1, \theta_1) \dots f(x_M, \theta_1)}{f(x_1, \theta_0) \dots f(x_M, \theta_0)} < A, \quad (4.8)$$

$$A \leq \frac{p_{1M}}{p_{0M}}. \quad (4.9)$$

where the p_{ij} is the probability of observing the length j sequence under hypothesis i . Note that M is a random variable since the observations are coming from a probability

distribution. At the M^{th} step we have crossed the threshold and declared H_1 . At this step consider $P_\theta(H_1|H_0) = P_\theta(x_1, \dots, x_M|H_0) = p_{0M}$. We will define this to be

$$P_\theta(H_1|H_0) = \alpha. \quad (4.10)$$

Similarly We will let

$$P_\theta(H_0|H_1) = \beta, \quad (4.11)$$

which makes $p_{1M} = 1 - \beta$, that is the probability of correctly declaring H_1 . We can then establish a key inequality by substituting into Eq. (4.9).

$$A \leq \frac{1 - \beta}{\alpha}. \quad (4.12)$$

A similar analysis of sequences that result in the declaration of H_0 will yield the following:

$$B \geq \frac{\beta}{1 - \alpha}. \quad (4.13)$$

Supposes that we have fixed an α and β corresponding to our two classical types of errors. In order for this test to achieve these error rates, the thresholds should be set to values $A = A(\alpha, \beta)$ and $B = B(\alpha, \beta)$ which are the proper functions of α, β respectively. We know from the previous inequality that:

$$\begin{aligned} A(\alpha, \beta) &\leq \frac{1 - \beta}{\alpha}, \\ B(\alpha, \beta) &\geq \frac{\beta}{1 - \alpha}. \end{aligned} \quad (4.14)$$

Because the assumed distributions only differ by a parameter, the ratio $\frac{p_{1m}}{p_{0m}}$ should be strictly monotonic with respect to increasing m . Suppose we set $A = A'$, where $A' > A(\alpha, \beta)$, then the test will terminate later at some $M' > M$ because the LRT was heading in the direction of the proper $A(\alpha, \beta)$ and monotonicity will cause it to cross A' eventually. This later thresholds should yield the same decision, but with a longer sequence. A similar argument applies to B , thus we can choose for A and B :

$$\begin{aligned} A &= \frac{1 - \beta}{\alpha}, \\ B &= \frac{\beta}{1 - \alpha}. \end{aligned} \quad (4.15)$$

These assignments of A, B can be shown (by explicit calculation, [16]) to not increase the required steps by a significant amount. The achieved error rates will in general not be equal to α, β though. For the values of A, B given by Eq.(4.15) let the effective error rates be given by α', β' , then from Eq.(4.14) we have have

$$\begin{aligned} \frac{1 - \beta'}{\alpha'} &\geq A = \frac{1 - \beta}{\alpha}, \\ \frac{\beta'}{1 - \alpha'} &\leq B = \frac{\beta}{1 - \alpha}, \end{aligned} \tag{4.16}$$

from which we can derive

$$\begin{aligned} \alpha' &\leq \frac{\alpha}{1 - \beta}, \\ \beta' &\leq \frac{\beta}{1 - \alpha}. \end{aligned} \tag{4.17}$$

Adding the equations of Eq.(4.17), we get the result

$$\alpha' + \beta' \leq \beta + \alpha, \tag{4.18}$$

which establishes a bound on the effective size power.

4.5 Types of tests

As we saw in section (4.1) hypothesis tests can come in a variety of flavors. The main classes are discrete versus continuous and composite versus simple hypothesis. For this problem we will primarily be concerned with discrete tests of a composite hypothesis. We will specifically focus on Bernoulli random variables $B(p)$ for the treatment of sequential tests as this is the case we will need to classify channels.

Let us initially consider a composite hypothesis formulated as:

$$\begin{aligned} H_0 : p &\leq p', \\ H_1 : p &> p'. \end{aligned} \tag{4.19}$$

p' is a soft-threshold, that is we deem p' the soft threshold because our test does not demand $p = p'$. In fact our test will only make better decisions for values of $p \ll p'$ or $p \gg p'$. For values of p that are close to p' , if our system declares the wrong hypothesis,

this should not be regarded as a serious error. Thus our hypothesis is equivalent to a test where we set hard thresholds, which are chosen to be far enough away from p' that making a mistake would be considered a serious error. Let p_0 and p_1 be these hard thresholds and define

$$\omega_r = \{p | p_1 < p\}, \quad (4.20)$$

be the region of rejection, that is the region where we declare H_1 . p_1 is chosen so that $p_1 > p'$. Similarly let

$$\omega_a = \{p | p < p_0\}, \quad (4.21)$$

be the region of acceptance, there region where we declare H_0 . The set of parameters

$$\omega_I = \{p | p_0 < p < p_1\}, \quad (4.22)$$

lie in the region of indifference. p_1 is a hard threshold in the sense that we expect that when the true parameter is equal to p_1 (that is when H_1 is true), we will make a mistake with probability β . Similarly when the true parameter is equal to p_0 (when H_0 is true), we will make a mistake with probability α . For all other values of the parameter in ω_r , α serves as a bound of the probability of making an type 1 error. The opposite is true for ω_a with β serving as the bound. In this test $(\alpha, 1 - \beta)$ together define the **size/power** of the test which are now both fixed. For a given p as the sequence length grow, the error probabilities are bounded by (α, β) . This establishes a bound for what to expect from the test performance. For a given size/power choice we can adjust p_0 and p_1 to try to minimize the expected sequence length $E_p(M)$. This however is not a trivial adjustment, and there is no “perfect minimum” as we will see when we examine the trade offs.

With these definition of the decision regions, we can now establish a relationship between simple hypotheses and composite hypotheses. Suppose that instead of Eq.(4.19) our hypothesis is simpler:

$$\begin{aligned} H_0 : p &= p_0, \\ H_1 : p &\neq p_0, \end{aligned} \quad (4.23)$$

so the region of accept has only a single value in it. That is $\omega_a = \{p_0\}$. This observation leads us to the distinction between simple and composite hypotheses in this context, namely the cardinality of the acceptance and rejection regions. The choice of hypothesis will have an impact on the fundamental curves that govern the behavior of our tests. In particular the operator curve (defined in section 4.1) will have a jump discontinuity in it when the hypothesis is simple. We can model simple hypothesis with a ϵ sized ω_a about p_0 , so that the operator curve is not discontinuous.

4.6 Fundamental Curves

We define the operator characteristic curve $L(p)$ as the probability of the test ending in a declaration of H_0 given a value of $p \in (0, 1)$. In section (4.1), we saw that $L(p) = E_\theta(\phi(X))$, where in this case our $\phi(x)$ is the SPRT. $L(p)$ is our fundamental tool for analyzing parameter choices. While it is possible to derive an expression for $L(p)$ from the integral in Eq. (4.4), Wald derives it from a simple continuity argument in [16]. In particular for the given composite hypothesis, we must have

$$L(0) = 1, \tag{4.24}$$

$$L(1) = 0,$$

$$L(p_0) = 1 - \alpha,$$

$$L(p_1) = \beta.$$

Our assumption of a Bernoulli random variable, plus the given composite hypothesis means that $L(p)$ admits a simple form, it is given by the parametric equations:

$$L(p) = \frac{\left(\frac{1-\beta}{\alpha}\right)^h - 1}{\left(\frac{1-\beta}{\alpha}\right)^h - \left(\frac{\beta}{1-\alpha}\right)^h}, \tag{4.25}$$

$$p = \frac{1 - \left(\frac{1-p_1}{1-p_0}\right)^h}{\left(\frac{p_1}{p_0}\right)^h - \left(\frac{1-p_1}{1-p_0}\right)^h}, \tag{4.26}$$

where $h \in (-\infty, \infty)$. There is a discontinuity at $h = 0$ which is due to the tests inability to gracefully handle the case where $p = p'$, but it's a simple point discontinuity and so can be safely ignored.

A second fundamental curve which can be considered a result of the parameter choices is $E_p(M)$, the expected number of samples required before a decision is made. It depends on the actual value of p . It has the form

$$E_p(M) = \frac{L(p)\log(\frac{\beta}{1-\alpha}) + (1-L(p))\log(\frac{1-\beta}{\alpha})}{p\log(\frac{p_1}{p_0}) + (1-p)\log(\frac{1-p_1}{1-p_0})}. \quad (4.27)$$

For given choices of p_0, p_1, α, β these curves can be drawn (see section 4.8).

4.7 Derivation of $L(p)$

Recall that we defined $L(p)$ as the probability of declaring H_0 when the test ends. There is a small subtlety in the phrasing of this definition that can be misleading. $L(p)$ is not the probability of declaring H_0 at an instant. That probability is dictated by the past observations and the derived thresholds and is determined by how close you are to the resulting H_0 threshold lines (defined in section 4.9). Instead $L(p)$ is the probability that the test will terminate with a declaration of H_0 . Note that even though there are 3 decision regions, the test can only ever terminate with one of two choices, H_0 or H_1 .

Consider the expression

$$\left[\frac{f(x, \theta_1)}{f(x, \theta_0)} \right]^{h(\theta)}, \quad (4.28)$$

where θ is a generic parameter of the distribution (not necessarily limited to the Bernoulli case).

For a given value of x , the quantity inside the $[\cdot]$ is fixed, and Eq.(4.28) is purely a function of θ . $h(\theta)$ is chosen such that the mean across x with respect to $f(x, \theta)$ equals one:

$$\sum_x \left[\frac{f(x, \theta_1)}{f(x, \theta_0)} \right]^{h(\theta)} f(x, \theta) = 1, \quad \forall \theta. \quad (4.29)$$

This choice of $h(\theta)$ ensures that

$$f^*(x, \theta) = \left[\frac{f(x, \theta_1)}{f(x, \theta_0)} \right]^{h(\theta)} f(x, \theta), \quad (4.30)$$

is a distribution on x for all allowable θ , it simply re-weights the individual assignments of $f(x, \theta)$. Fix a θ and let H be the hypothesis that $f(x, \theta)$ is the true distribution on x . Let H^* be the hypothesis that $f^*(x, \theta)$ is the true hypothesis, where $f^*(x, \theta)$ comes

from Eq. (4.30). If we consider the LRT of a sequence of observations x_1, \dots, x_m , we have

$$B^{h(\theta)} < \left[\frac{f^*(x_1, \theta) \dots f^*(x_m, \theta)}{f(x_1, \theta) \dots f(x_m, \theta)} \right] = \left[\frac{f(x_1, \theta_1) \dots f(x_m, \theta_1)}{f(x_1, \theta_0) \dots f(x_m, \theta_0)} \right]^{h(\theta)} < A^{h(\theta)}. \quad (4.31)$$

It is shown in the appendix of [16] that for a small set of assumptions on $f(x, \theta)$, $h(\theta)$ takes on only one non-zero value if it is to satisfy Eq. (4.29) for each θ . We can treat two cases $h(\theta) > 0$ and $h(\theta) < 0$. They yield the same result however, namely for given set of observations x_1, \dots, x_m , if the left LRT of Eq. (4.31) declares H , then the right LRT will declare H_0 and vice versa. Since the tests yield the same results for a given set of observations, it follows that the probability of the test ending in the declaration of H_0 is the probability that the left LRT declares H .

We saw in section (4.4) that for any LRT, and choice of α, β , we can take the upper threshold to be $\frac{1-\beta}{\alpha}$ and the lower threshold to be $\frac{\beta}{1-\alpha}$. While this is not exact, taking these thresholds ensures that our $\alpha_{actual} < \alpha$ and $\beta_{actual} < \beta$. In this case, we can use similar results to approximate the bounds. Let

$$\begin{aligned} A^{h(\theta)} &\sim \frac{1-\beta'}{\alpha'}, \\ B^{h(\theta)} &\sim \frac{\beta'}{1-\alpha'}, \end{aligned} \quad (4.32)$$

where α', β' are the respective probabilities of miss and false alarm. From these equations we can solve for α' and get

$$\alpha' = \frac{1 - B^{h(\theta)}}{A^{h(\theta)} - B^{h(\theta)}}. \quad (4.33)$$

Given that α' is the probability of declaring H_1 when the parameter is in the accept region, we have $\alpha' = 1 - L(\theta)$, which yields

$$L(\theta) = \frac{A^{h(\theta)} - 1}{A^{h(\theta)} - B^{h(\theta)}}. \quad (4.34)$$

Now in the Bernoulli case Eq.(4.29) yields,

$$p \left(\frac{p_1}{p_0} \right)^h + (1-p) \left(\frac{1-p_1}{1-p_0} \right)^h = 1, \quad (4.35)$$

where h is the nonzero value of $h(p)$ for that specific p . Solving Eq.(4.35) for p and performing the proper substitutions into Eq.(4.34), yields Eq.(4.25) and Eq.(4.26).

4.8 Curve plots

In this section the following parameter choices were made. Here we will introduce a parameter δ which is an offset from p' . Then the p_i for $i \in \{0, 1\}$ can be computed as

$$p_1 = p' + \delta, \quad (4.36)$$

$$p_0 = p' - \delta.$$

This has the effect of putting p' in the middle of the indifference region. The δ parameter will be explored in more detail later. For this sample curve the parameters are:

$$\alpha = 0.01, \quad (4.37)$$

$$\beta = 0.05,$$

$$p' = 0.20,$$

$$\delta = 0.15.$$




Figure 4.1: Operator Curve

Figure (4.1) contains two truncated curves the theoretical curve which was calculated from Eq.(4.25) and Eq.(4.26), and a simulated curve that comes from running the

sequential experiments and then collecting the actual $p(H_0)$. They were truncated to show the interesting detail as the edge behavior is asymptotic. Note that the theoretical curve passes through the expected points.

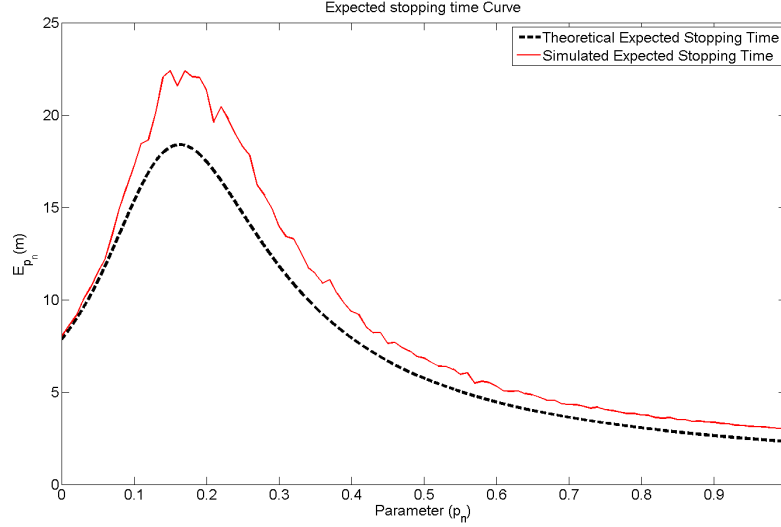


Figure 4.2: Expected sample length curve

Figure (4.2) shows the expected sequence length. As you can see the closer the true parameter is to the soft threshold the more samples we need to verify the hypothesis. This confirms the intuition that the closer the parameter is to the threshold, the harder it is to tell what side the parameter is on. This is an important result since p' will shape many of our design decisions.

4.9 Threshold lines and Simulation

Applying some algebraic manipulation and logarithms to Eq.(4.7) and using the Bernoulli distribution as our model, we find that $d_m = \sum_{m=0}^M x_m$ is a sufficient statistic for the proposed test (see [15] for a discussion of sufficiency). Here d_m is the number of ones observed after m samples. We can plot this quantity as a function of sample index and compare it to two threshold lines that come from the sufficiency derivation. If d_m lies between the two threshold lines $L_0(m)$ and $L_1(m)$ then we continue sampling. We declare hypothesis H_i when we cross $L_i(m), i \in \{0, 1\}$. The formulas for the threshold

lines are given as

$$\begin{aligned} L_0(m) &= \frac{\log(\frac{\beta}{1-\alpha})}{\log(\frac{p_1}{p_0}) - \log(\frac{1-p_1}{1-p_0})} + m \frac{\log(\frac{1-p_0}{1-p_1})}{\log(\frac{p_1}{p_0}) - \log(\frac{1-p_1}{1-p_0})}, \\ L_1(m) &= \frac{\log(\frac{1-\beta}{\alpha})}{\log(\frac{p_1}{p_0}) - \log(\frac{1-p_1}{1-p_0})} + m \frac{\log(\frac{1-p_0}{1-p_1})}{\log(\frac{p_1}{p_0}) - \log(\frac{1-p_1}{1-p_0})}. \end{aligned} \quad (4.38)$$

A sample run looks like figure (4.3). Note that the lines are parallel (equal slope) and only differ by their intercept. This observation was used in the proof of eventual termination of the SPRT.

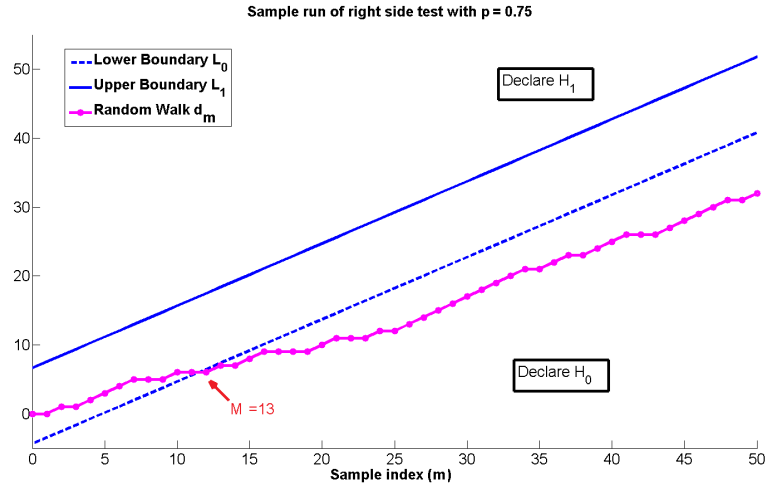


Figure 4.3: Example of the Test

To generate the simulated curves in 4.8, we sample a Bernoulli RV with parameter $p \in [0, 1]$ until it crosses one of these threshold lines. Performing this process over thousands of trails for all $p \in [0, 1]$, we can analyze the resulting declarations and sequence lengths, this is examined in section (6.1).

4.10 Random walks between two thresholds

An alternate way to think about the SPRT is to consider it as a random walk between two moving thresholds as in figure (4.4). This can be seen as projecting the graph in figure(4.3) onto the y axis. The sufficient static $d_m = \sum_{m=0}^M x_m$ is by definition a random walk as it is the sum of an increasing series of IID Bernoulli random variables [12]. As m increases the thresholds move to the left (this is convention, there is no

specific reason for left vs right). If the parameter p is low then d_m will get overtaken by the lower threshold L_0 , causing the test to decide H_0 . On the other hand, if the p is high, then d_m will over take L_1 , causing the test to decide H_1 . One might try to derive $E_p(M)$ directly from Wald's equality $E[d_m] = E[X]E[M]$. In [15] it is shown that in the case of a simple random Wald's equality yeilds:

$$P\{\exists M \text{ s.t. } d_M \geq T\} = \left(\frac{p}{1-p}\right)^T.$$

The problem with this equation is that it assumes a fixed T . While this method can not be used directly to derive $E_p(M)$, it does suggest techniques like considering the ratio of the expected sum to the expected value of the random variable, i.e., $\frac{E[d_m]}{E[X]} = E[M]$.

With this formulation it's clear that the SPRT has a Markovian property since at any m , the probability of deciding H_0 or H_1 is purely a function of the distance from one of the threshold lines to d_m . Note that if the p is closer to 0 or 1, this test will complete in fewer samples because both threshold lines move at the same rate (slopes are equal). This can also be seen from $E_p(M)$ curves of section(4.8).




Figure 4.4: Random Walk between two thresholds

4.11 Computing $\beta_{effective}$

By definition $\beta = P\{\text{Declare } H_0 | H_1 \text{ is true}\}$. $L(p)$ is by definition $P\{\text{Declare } H_0 | p\}$. Since H_1 is true when $p > p'$, we can compute the $\beta_{effective}$ by integrating $L(p)$ over the region where H_1 is true. In the most general case:

$$\beta_{effective} = \int_{p'}^1 L(p) f(p) dp, \quad (4.39)$$

For simplicity we can take $f(p) = \frac{1}{\mu(S)} = \frac{1}{1-p'}$ where $S = (p', 1)$, that assuming a uniform distribution on $p \in (p', 1)$. Clearly $S = (p', 1) \supset \omega_r = (p_1, 1)$ so we can establish a relationship between $\beta_{effective}$ and β like so:

$$\beta_{effective} = \frac{1}{\mu(S)} \int_{p'}^1 L(p) dp, \quad (4.40)$$

$$= \frac{1}{\mu(S)} \left[\int_{p'}^{p_1} L(p) dp + \int_{p_1}^1 L(p) dp \right]. \quad (4.41)$$

Since $L(p)$ is monotonically decreasing for $p > p'$ we have:

$$\beta \mu(\omega_r) \geq \int_{p_1}^1 L(p) dp, \quad (4.42)$$

$$\beta \frac{\mu(\omega_r)}{\mu(S)} + \frac{1}{\mu(S)} \int_{p'}^{p_1} L(p) dp \geq \beta_{effective}, \quad (4.43)$$

$$\beta \geq \beta \frac{\mu(\omega_r)}{\mu(S)} \geq \beta_{effective} - \frac{1}{\mu(S)} \int_{p'}^{p_1} L(p) dp, \quad (4.44)$$

The first inequality of (4.44) follows because $1 \geq \frac{\mu(\omega_r)}{\mu(S)}$. Clearly as $p_1 \rightarrow p'$, $\beta_{effective}$ grows closer to β which becomes the true false alarm probability. (4.44) shows how choosing looser (or tighter) p_1 will effect the effective false alarm rate, $\beta_{effective}$. As we saw in section (4.8), the closer p_1 is to p' the higher the $E_p(n)$ becomes. For that cost we do gain a reduction in false alarms.




Figure 4.5: The over shoot of $\beta_{effective}$

We can also examine this intuition from the graph of $L(p)$ as in figure 4.5. Since $L(p_1) = \beta$, as $p_1 \rightarrow p'$, this forces the curve to pass through the value β earlier because

we require $L(p)$ to be smooth. This shrinks the effective area of the shaded region thus cutting out significant area under the $L(p)$ curve over the region $(p', 1)$ and thus lower the value of the integral in Eq. (4.40). A similar analysis can be done for α and p_0 .

4.12 $\alpha_{effective}, \beta_{effective}$ and δ

If we consider Eq. (4.41) and adjust the equations to reflect the $\beta_{effective}$'s dependence on δ , we derive

$$\beta_{effective}(\delta) = \int_{p'}^{p'+\delta} L(p)dp + \int_{p'+\delta}^1 L(p)dp, \quad (4.45)$$

$$\leq \int_{p'}^{p'+\delta} L(p)dp + L(p_1)\mu(\omega_r), \quad (4.46)$$

$$\leq L(p')\delta + L(p'+\delta)(1 - (p'+\delta)). \quad (4.47)$$

When δ is really small the fixed point $L(p_1) = L(p'+\delta)$ on the operator curve is very close to p' , so (4.47) is dominated by the second term. Monotonicity of $L(p)$ in the region $p > p'$ implies that $L(p') > L(p'+\delta)$. Clearly as $\delta \rightarrow 0 \implies \beta_{effective} \searrow$ as was noted in section (4.11).

We can see this effect from figure (4.6). This graph is the $\beta_{effective}$ integral given in Eq. (4.40) numerically integrated as $\delta : 0.01 \rightarrow 0.09$. The graph's growth is in accordance with intuition from the treatment of the bound.




Figure 4.6: δ vs β

A sample $E_p(N)$ curve is shown in figure (4.7) with $\delta = 0.09$, $\beta_{effective}(\delta) = 0.0437$, and $\max(E_p(M)) = 26.3996$.




Figure 4.7: Large delta (big indifference region) expected sequence length

At the same time we can bound $\alpha_{effective}$ by:

$$\alpha_{effective}(\delta) = \int_{p'-\delta}^{p'} L(p)dp + \int_0^{p'-\delta} L(p)dp, \quad (4.48)$$

$$\leq \int_{p'}^{p'-\delta} L(p)dp + L(p_1)\mu(\omega_a), \quad (4.49)$$

$$\leq L(p')\delta + L(p' - \delta)(p' - \delta), \quad (4.50)$$

which again has the same relationship with δ namely $\delta \rightarrow 0 \implies \alpha_{effective} \searrow$. As we saw in 4.11, this reduction in error comes at the cost of $\max(E_p(M))$.

While adjusting all the other parameters has an effect on both $\max(E_p(M))$ as well as $(\alpha_{effective}, \beta_{effective})$, these quantities are most sensitive to changes in δ , because changes in δ affect the declaration regions the most. $\max(E_p(M))$ in particular primarily depends on the width of the indifference region, which is 2δ . This maximum occurs at the center of the indifference region, which was set to p' for all the previous analysis.

The fundamental trade off identified here is an exchange of the expected sequence length for the probability of error. However, this is not the only exchange we can make. We can introduce a new parameter ζ which represent the split of the indifference width.

We can compute the values of p_1 and p_0 as $p_1 = p' + 2\zeta\delta$, $p_0 = p' - 2(1 - \zeta)\delta$. This moves the center of the indifference region and the position at which $\max(E_p(M))$ occurs. This has the effect of trading $\beta_{effective}$ for $\alpha_{effective}$ as the integration regions get adjusted to:

$$\begin{aligned}\alpha_{effective}(\delta) &= \int_{p'-2(1-\zeta)\delta}^{p'} L(p)dp + \int_0^{p'-2(1-\zeta)\delta} L(p)dp, \\ \beta_{effective}(\delta) &= \int_{p'}^{p'+2\zeta\delta} L(p)dp + \int_{p'+2\zeta\delta}^1 L(p)dp.\end{aligned}\tag{4.51}$$

This gives us another mechanism for adjusting the SPRT to have low $E_p(M)$ over the region of where the parameter p highest probability. The effects of this ζ term can be seen in figure (4.8).

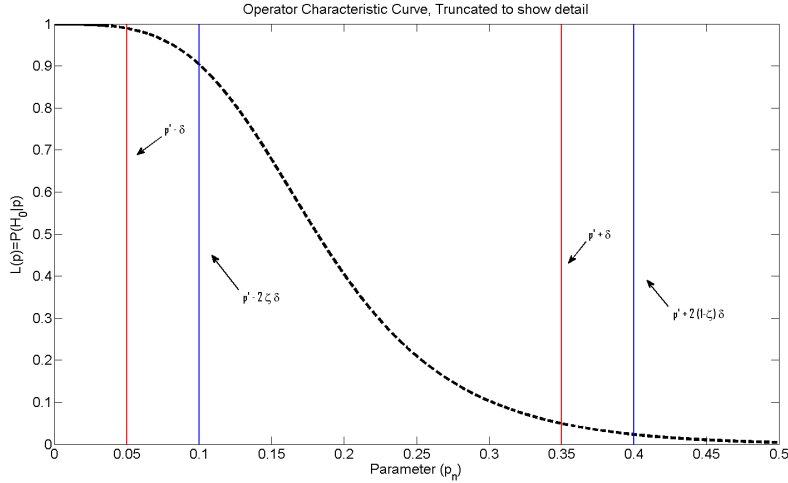


Figure 4.8: Shift of the indifference region boundaries due to ζ

4.13 Multiple SPRTs

Hypothesis like Eq.(4.19) partition the parameter space into two disjoint regions, $R_0 = \{p|p < p'\}$ and $R_1 = \{p|p' < p\}$. If we examine the variance of a Bernoulli random variable (see figure 4.9), we notice that it is symmetric about the parameter $p = \frac{1}{2}$. Because of this symmetry, for properly chosen parameters, we can run two tests simultaneously without affecting the expected sequence length $E_p(M)$ or the effective error rates $(\alpha_{effective}, \beta_{effective})$ significantly. By running 2 tests against the same

set of observations we will partition the parameter space into 3 intervals using two thresholds p'_L, p'_R ($p'_L < p'_R$). This will partition the parameter space into three regions $R_0 = \{p | p < p'_L\}$, $R_1 = \{p | p'_L < p < p'_R\}$ and $R_2 = \{p | p'_R < p\}$. We expect to oper-




Figure 4.9: Variance of a Bernoulli Random Variable

ate in a regime where samples are limited but computation is unconstrained (limited sampling time, but infinite power). Under this assumption, we can get a better classification granularity by computing multiple SPRTs with the same set of samples. As we saw in section (4.2) we can use an elimination process to choose 1 hypothesis from many. For this discussion we will consider the case of $\Gamma = 3$. It is possible to use larger Γ but the error rates become much more difficult to analyze. The three hypotheses corresponds to the three regions or parameter space is partitioned into:

$$H_0 : p < p'_L, \tag{4.52}$$

$$H_1 : p'_L < p < p'_R,$$

$$H_2 : p'_R < p.$$

We can decide between these three by using the elimination process against two equivalent sets of binary hypothesis:

$$\begin{array}{ll}
 \text{LEFT} & \text{RIGHT} \\
 H_1 : p < p'_L & H_1 : p > p'_R, \\
 H_0 : p > p'_L & H_0 : p < p'_R.
 \end{array} \tag{4.53}$$

Each of these tests will generate a set of threshold lines $L_{0,L}, L_{1,L}$ and $L_{0,R}, L_{1,R}$. As was observed in [17] we can run these tests simultaneously by simply plotting both sets of test lines, and using a set of combining rules to decide what the final declaration of the test will be. A sample run of two would look like figure (4.10). The rule for combining

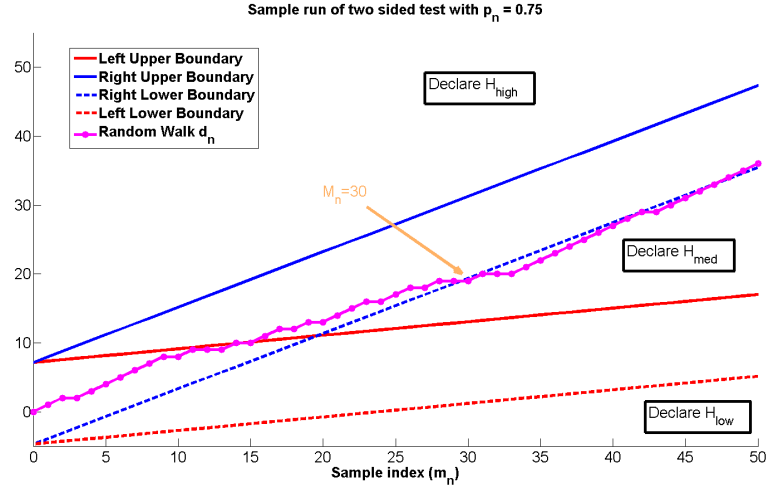


Figure 4.10: Sample run of 3-ary test with 2 sets of threshold lines

the decisions from the two tests into a final declaration decides in the following manner:

- Left declares H_0 and Right declares H_1 , declare $H_0 : p < p'_L$
- Both declare H_1 , declare $H_1 : p'_L < p < p'_R$
- Right declares H_0 and Left declares H_1 , declare $H_2 : p'_R < p$

We can see that from figure (4.10), that this decision rule merely picks a hypothesis based on what set of threshold lines d_m crossed. If d_m crosses one of the outer threshold line, a declaration of H_0 or H_2 is made depending on which line was crossed. If d_m

ends up in the region bounded by the inner 2 threshold lines, H_1 is declared. It is argued in [17] that contradicting declarations cannot occur because that would violate the convergence of the single test.

4.14 $L(p)$ and $E_p(M)$ for two-tests

If we examine each test in section (4.13) individually, it will have an $E_p(M)$ similar to figure (4.2) for a single test. Note how this curve is unimodal with a maximum near p' , the soft threshold for the two hypothesis test. We see that as p moves away from p' in either direction, $E_p(M)$ decreases monotonically. If we now consider the two tests together there will be two soft thresholds p'_L and p'_R . An arbitrary p will tend to be closer to one of the soft thresholds than the other (the case where p is centered between the two does not cause any discontinuities). Because of the monotonicity of the individual $E_p(M)$, one test will always finish before the other. Since we can only make a decision when we have answers from both tests, see that $E_{2,p}(M)$, the two-tests expected sequence length, can be computed as $E_{2,p}(M) = \max(E_{L,p}(M), E_{R,p}(M))$, where $(E_{L,p}(M), E_{R,p}(M))$ are the individual left and right $E_p(M)$ curves respectively.

In general $(E_{L,p}(M), E_{R,p}(M))$ will not be symmetric, so the maximum function will always pick one side. We can however simplify the test and calculations by picking $p'_R = 1 - p'_L$, and $\delta_L = \delta_R$. It is assumed that $p'_L < \frac{1}{2}$, otherwise the decision regions will be degenerate. In practice, $0.3 < p'_L$ yields poor performance because the indifference regions are too close together. For a proper choice of p'_L and p'_R , the resulting $E_{2,p}(M)$ curve will look like figure (4.11)

For $\Gamma = 3$ the original operator curve is not very meaningful because purely paying attention to declarations of H_0 ignores the 3rd hypothesis. The errors of this test become difficult to analyze because there are now two kinds of type I error for each hypothesis. An operator curve might not be feasible in this case because some of the points that the curve are required to pass through might contradict, e.g., incorrectly declaring H_1 instead of H_0 might have a different probability than incorrectly declaring H_1 instead of H_2 .

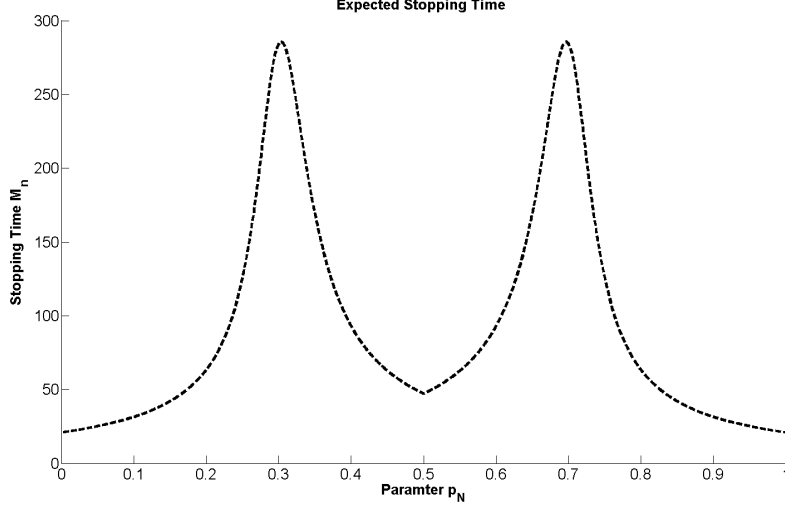


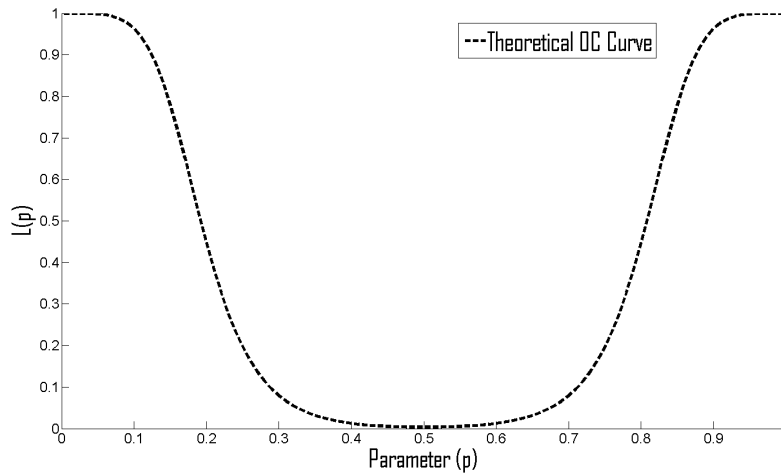
Figure 4.11: Expected Sequence length, $E_{2,p}(M)$, for a combined two-test

It is possible however to reduce the problem to the two hypothesis case by simply “gluing” together the two outer hypothesis, for the purposes of error analysis. The test will still return 1 of 3 answers, but we will treat the errors that come from declaring H_0, H_2 as the same. If we stick to symmetrically chosen parameters then the error analysis actually simplifies quite well. If we take as our hypothesis,

$$H'_0 : p > p'_L \text{ or } p'_R < p, \quad (4.54)$$

$$H'_1 : p'_L < p < p'_R,$$

and choose our parameters as before, $p'_R = 1 - p'_L$, and $\delta_L = \delta_R$, then the probability of declaring H_0 is simply the probability that either the left or the right declares H_0 . Since these tests run independently of each other, $L_2(p) = L_L(p) + L_R(p) = L_L(p) + L_L(1 - p)$, where $L_2(p)$ is the two-test operator curve, and $(L_L(p), L_R(p))$ are the left and right individual operator curves respectively. The last equality follows because of the symmetry. The two-test operator curve can be seen in figure (4.12), it has the expected symmetry about $p = \frac{1}{2}$. From the operator curve we can again compute $(\alpha_{effective}, \beta_{effective})$ by integrating the operator curve over the accept and reject regions as was done in section (4.11). In this case the do not reject region is $\omega_a = \{p | p > p'_L\} \cup \{p | p'_R < p\}$ and the reject region is $\omega_r = \{p | p'_L < p < p'_R\}$.

Figure 4.12: The $L_2(p)$ OC curve

4.15 Parameter sensitivity

For a given size/power $(\alpha, 1 - \beta)$ of test we can adjust the accept/reject regions boundaries (that is we can choose different p_1, p_0), which has an effect of on the slope and intercept of L_0, L_1 .

If we consider δ where $\delta = \frac{p_1 - p_0}{2}$, we can make δ larger by choosing p_1, p_0 which are separated by a further distance. This will result in smaller accept and reject regions. These smaller accept/reject regions cause two effects. $E_p(M)$ becomes lower because we are making decisions on parameters that are much further away from the threshold. At the same time, however, the indifference region gets bigger. Even though our false alarm bound is fixed, our effective false alarm / miss probability are going to suffer because the test is now indifferent to larger region about $\frac{p_1 - p_0}{2}$. On the other hand we can shrink δ to a very tight region around $\frac{p_1 - p_0}{2}$, this will give us very good effective error probabilities, but will come at the cost of much higher $E_p(M)$. In figure (4.13) and figure (4.14) we can see the effect of widening δ . Note that as $L(p)$ becomes shaped more like a step function, $\max(E_p(M))$ grows unboundedly.

For a given (α, β) pair we can compare bounds on the number of samples required for the UMP to achieve these error rates against the $E_p(M)$. In [16] a bound for the




Figure 4.13: Operator curves with fixed p' and varying δ




Figure 4.14: Expected sequence length with fixed p' and varying δ

UMP's required samples to achieve (α, β) is given as:

$$M \approx \frac{(Q^{-1}(\beta) - Q^{-1}(1 - \alpha))^2}{(p_0 - p_1)^2}, \quad (4.55)$$

where Q is the Q -function of a standard normal random variable. The key observation is that this value has no p dependence and thus is constant across the range of p . In contrast, the $E_p(M)$ can be orders of magnitude below $E_p(M)$ which occurs at the center of the indifference region. If we know something about how the space of possible p 's is distributed, for a fixed (α, β) , we can tune the SPRT to achieve a lower sequence lengths in the regions of interest. That is there is a choice of p_1, p_0, α, β which positions the low portion of the $E_p(M)$ in the optimal portion of the parameter space.

Chapter 5

The Cost Classifier

In section (4) we saw that the SPRT classifies the channel with the lowest expected sequence length. Tuning the SPRT amounts to choosing the bins we place our channels in (p'_L, p'_R mark the boundaries), and choosing a balance between $(\alpha_{effective}, \beta_{effective})$ and the $E_p(M)$. We have simplified the trade off to a choice of δ and ζ . As noted in section (4.10), the Markovian property of this test implies that the expected sequence length after an arbitrary number of steps is purely dependent on the probability of occupancy p and the current proximity to the threshold lines.

5.1 Cost Definition and Greedy Allocation

Based on the observations of section(5) we propose a cost metric to guide our measurement distribution across all channels. This metric rewards movement towards threshold lines and penalizes excessive sequence lengths. The cost of channel n is then defined to be:

$$C(m_n, n) = \begin{cases} e_{n, m_n} + (\lambda * m_n) & : \text{unclassified,} \\ \infty & : \text{classified,} \end{cases} \quad (5.1)$$

where e_{n, m_n} is the distance to the nearest outer threshold line, $n \in [0, |N|]$ is the channel index, $m_n \in [0, M_n]$ is the number of samples given to channel n (it's local clock). Here M_n is the stopping time for channel n . Each channel has it's own m_n , so to simplify the notation, we'll simply use m when m_n is paired with n , e.g. $C(m, n)$. λ is the measurement penalty coefficient, a constant that does not depend on n .

As the search progresses the classifier will dwell on some channels that do not complete thier classification because these channels will take too long according to the cost. We index these channels by n' . Each of these channels has a dwelling time denoted

by $M'_{n'}$. Channels indexed by n'' do complete their classification with a stopping time given by $M_{n''}$. We define a global time index j as $j = \sum_{n''} M_{n''} + \sum_{n'} M'_{n'}$. Note that $j \in [0, L]$ (recall that L is our measurement limit) and the sum $\sum_{n'} M'_{n'}$ will contain the channel currently under observation. Once channel is classified, it will belong to the sum $\sum_{n''} M_{n''}$.

The greedy selection rule for each measurement choice is then $\min_n (C(m, n))$. At each tick of the global clock j we pick the channel with the smallest cost. As the random walk $d_{n,m} = \sum_m x_{n,m}$, the sum of the m samples from channel n , tends towards an outer threshold line $e_{n,m}$ will decrease. If a channel n receives enough measurements to complete classification there is no reason to measure it any more, hence setting the cost to ∞ means the selection rule will never come back to this channel.

The quantity $e_{n,m}$ is a measure of how close the test for channel n is to completion. If the sum $d_{n,m}$ moves away from these outer threshold lines towards the center of either pair of threshold lines, the $\lambda * m$ term will drive the cost of this channel up. It is assumed that all channels initially start with the same cost which is due to our assumption that initially all channel parameters that are unknown and uninformed. If a channel's cost is driven up past that of an unexplored channel, the selection rule will ignore it until all other channels have had a chance to generate a cost reduction. In a low measurement regime, channels with cost that are above the cost of an unexplored channel will probably never be revisited. If there are several choices for $\min_n (C(m, n))$, we pick one randomly.

There is a bias in the allocation because the distance metric favors channels that are at the edges of the interval $(0, 1)$. Because of the way our distance metric is defined, we favor H_0 of the Eq. (4.54). We picked this metric of proximity to completion because we observed in section (4.13) that channels with parameters p at the edges have shorter expected sequence length. If we prioritize these channels over ones that are in near the soft thresholds (p'_L, p'_R) we will spend fewer samples characterizing individual channels at the beginning of the search. If we have enough samples to measure all the channels $n \in [0, |N|]$ then we will eventually get to the channels that will take more samples to characterize. However, if we do not have enough samples, by prioritizing channels that

finish in fewer samples we maximize the number of channels we can classify.

5.2 Distance metrics

There are at least three different ways to compute $e_{n,m}$, however each of them results in essentially the same allocation process. The core difference between them is what units $e_{n,m}$ is measured in, and what values of λ are allowable due to the difference in units. The 3 different methods are as follows:

$$\begin{aligned} e_{n,m}^1 &= \min_h (|d_{n,m} - L_h(m)|), \\ e_{n,m}^2 &= \min_h (\text{dist}((m, d_{n,m}), L_h(m))), \\ e_{n,m}^3 &= \min_h (\lceil M'_{n,m,h} - m \rceil) \end{aligned} \tag{5.2}$$

The simplest to compute is $e_{n,m}^1$, it is just the vertical distance to the correspond threshold line. The minimization is over the hypothesis index $h \in \{0, 1\}$. This value is unitless. For $e_{n,m}^2$ the $\text{dist}()$ function is simply the Euclidean distance. The quantity $e_{n,m}^2$ is the distance to the nearest point on the threshold line, again minimized over h . The distance metric $e_{n,m}^1$ is vertical projection of $e_{n,m}^2$. Clearly as $e_{n,m}^2$ grows so does $e_{n,m}^1$. This quantity is also unit-less.

$e_{n,m}^3$ is a little more difficult to compute, the quantity $M'_{n,m,h}$ is the ideal stopping time from current position. It is the number of steps required to cross threshold line h assuming every sample after the current one was perfect. For example, if we were trying to cross the lower bound ($h = 0$) and were already used m samples, $N'_{m,0}$ would equal the number of zeros required to cross L_{lower} . All three distances are visualized in figure (5.1). For the case of the lower threshold line the ideal continuation would consistently add zero, whereas in the upper threshold line the ideal continuation would consistently add one.

In figure (5.1) we have a sample path with three ones and five zeros. This path is tending towards the lower threshold line. In this case $e_{n,m}^3$ will measure the steps to complete with respect to the lower threshold line. In this case $e_{n,m}^3$ is just the horizontal




Figure 5.1: Graphical representation of outer threshold distance calculation

projection of $e_{n,m}^2$. Again we have a similar relationship between these two distance metrics, $e_{n,m}^2$ decreasing $\implies e_{n,m}^3$ decreasing. Since each of these distance calculations obey the same growth trends, using any one of them as the $e_{n,m}$ term will yield the same hopping pattern, however that actual numerical value of $C(m, n)$ will be different. An equivalent relationship can be shown for the upper threshold lines.

Although all three metrics yield the same hopping strategy, there is one key difference between them. Metric 3 is measured in units of steps to complete, which gives a very intuitive interpretation to λ fractions of a step. If we choose lambda as some fraction $\frac{1}{z}$ then a test must take z steps to offset the cost reduction of 1 step towards the threshold lines. For our implementation we choose $e_{n,m} = e_{n,m}^3$.

5.3 Early mistakes

When using the two-test case, the outer threshold lines diverge. From section (4.14) we know that the two-test will complete in $E_{2,p}(M)$ samples. For the two-test case, this divergence self imposes a measurement penalty for channels that do not move toward an outer threshold line soon enough. Figure (5.1) shows that the divergence might cause early mistakes to be penalized unfairly. If channel n with parameter p_n gets an disagreement between samples early on, the jumping rule will skip over it before it can

generate enough cost reduction to offset the divergence cost penalties even if $\lambda = 0$. For example, if the first two samples of a channel i were $x_{n,1} = 0, x_{n,2} = 1$, then this channel has essentially not changed position, however it has consumed two samples, and thus moved forward two steps in figure (5.1). The value of $e_{n,m}$ has increased, however we still do not know much about this channel.

This effect is caused by the usage of the two-test, if we were using a single test then the parallel threshold hold lines would not exhibit this problem. To compensate for this early mistake sensitivity we can take one of two strategies. We can introduce a term to cost metric to prevent early jump. The cost function them becomes:

$$cost_m(c_n) = \begin{cases} e_{n,m} + (\lambda * m) + F(m) & : \text{ unclassified} \\ \infty & : \text{ classified,} \end{cases} \quad (5.3)$$

where $F(m)$ is a forgiveness factor defined as

$$F(m) = \begin{cases} -\varphi * m & m < m_f \\ 0 & m > m_f, \end{cases} \quad (5.4)$$

where the coefficient φ is chosen to offset the increase in cost due to the divergence until the m_f^{th} step. At this stage if the channel n has not generated enough cost reduction on it's own, then a jump is justified. There is no exact empirical rule for determining the forgiveness step m_f . As an alternative, for a given δ we can use the shifting capability of ζ to make the outer threshold lines closer in slope. We do this by shifting the decision regions closer together so that the slope terms are closer in value.

Since we assume that the channel parameters p_n are drawn from a uniform distribution $U[0, 1]$, we can use a simple argument to gauge what m_f should be. Consider the left side threshold, the probability of finding a channel below a threshold p'_L is simply the width of the interval $\mu([0, p'_L]) = p'_L$. Similarly the right side probability is $\mu([p'_R, 1]) = p'_R$. If we choose (p'_L, p'_R) as in section (4.14) then we have $P(p_n \in \omega_a) = 2p'_L$

If we treat the inspection of a channel as a simple Bernoulli trial, $B_{Inspection}(p_n) =$

$P(p_n \in \omega_a)$, then the number of channels we must inspect before we find one below the threshold is geometrically distributed with parameter $2p'_L$. Thus a lower bound on the expected number of channels we need to inspect before we find one that is below our threshold is $\frac{1}{2*p'_L}$. This simplification can be used as a guiding principle for the choice of m_f . The ideal number of channels inspected before a classification should not be much greater than $\frac{1}{2p'_L}$.

The principal argument against using the $F(m)$ term is that it offsets the effect of the single channel error rates on the cost function. It becomes harder to determine if we should have legitimately passed on a channel. On the other hand, when we adjust the ζ shift, the changes to the cost function are straight forward as the cost function definition remains unchanged.

5.4 Conditions for Optimality of the greedy approach

The selection rule is deemed greedy because at each step it seeks out the lowest cost choice without consideration of possible future cost reduction. For our model we assume that the parameters p_n for each channel n are constant during the search period. Let M_n be the stopping index for the n^{th} channel. In section (4.10) we noted that the $E_{p_n}(M_n)$ was only a function of the distance to the threshold line and parameter p_n . This property will be used to show the optimality of the greedy approach is tied to our choice of λ .

5.4.1 When is greedy approach optimal for our search

One of the requirements for the greedy approach to be optimal is that our test must possess the greedy-choice property, that is the greedy approach should be able to pick the globally optimal solution. To possess this property our model would need two qualities:

- Every channel with increases in cost will not decrease in cost for the duration of the search (monotonic increase).
- Every channel with decreasing in cost will continue to have decreasing until completion.

these two properties are necessary so that every choice of the lowest cost channel can not be replaced by a channel that at the current instant has a higher cost, but over the length of the sequence costs less.

5.4.2 Why greedy is not always optimal

Using the definition of greedy-choice property given in [18], we see that our allocation problem does not have this property in general. Indeed there may be cases where the greedy choice does not pick a globally optimal solution. Consider a simple example where we get a single inconsistent sample very early in our sampling from a channel n that with $p_n < p'_L$. Let $x_{n,1} = 0$, $x_{n,2} = 1$, $x_{n,3} = 0$, $x_{n,4} = 0$ be the first four samples of channel n . In this case the current estimate $\hat{p}_n = \frac{1}{4}$. Since channel n had an early disagreement, it's cost may go up if λ is high or if the outer threshold lines separate quickly. Channel n 's random walk returned to it's original unmeasured trajectory but consumed two samples. Because of this sequence of samples, channel n 's cost will be penalized as compared channels which have not yet taken any samples. In the large N regime, there will always be an unmeasured channel to pick over channel n .

In section (5.3) we have lower bounded the number of channels we need to inspect before finding a desired one by $\frac{1}{2p'_L}$. Suppose $p'_L = 0.2$, we can expect to go through ≈ 2.5 channels before we find another channel $n' \neq n$ with parameter $p_{n'} < p'_L$. If we give channel n two more samples, and get the result $x_{n,1} = 0$, $x_{n,2} = 1$, $x_{n,3} = 0$, $x_{n,4} = 0$, $x_{n,5} = 0$, $x_{n,6} = 0$ then $\hat{p}_n = \frac{1}{6} \approx 0.166\ldots$. The sufficient statistic $d_{n,6}$ moved significantly closer to the threshold line.

The most aggressive case of the two-test will jump to another channel if the first two samples don't agree. This choice would explore maximally but pass on many channels that satisfied the criteria $p < p'_L$. If we skip channel n and we spend at least two samples on channels we are not dwelling on, we would expect to spend $\approx 4 - 6$ samples before we found a channel that generates enough cost reduction to dwell on. In contrast had we stayed on channel n , we would have only spent two more samples and would have gotten a enough cost reduction to justify dwelling on channel n for longer. Here the greedy choice failed to be optimal because our cost function did not properly price

the channel.

5.4.3 Lower bound on the cost to characterize a single channel

As we saw in (5.3) the lower bound on channels we have to pass through before we find one worth expending samples on is $\frac{1}{2p'_L}$. The ideal lower bound on cost to characterize any channel is given as $C(n) = \sum_{m=0}^{M_n} C(n, m)$, where $m = 0$ is the initial cost. For each channel we pass on, the cost spent on them is $C'(n) = \sum_{m=0}^{M'_n} C(n, m)$ where M'_n is the number of samples spent (without completion). A lower bound to the cost required to classify channel n is then

$$C_{min}(n) = \sum_{n'} C'(n') + C(n) = \lfloor (\frac{1}{2p'_L}) C'(n) + C(n) \rfloor, \quad (5.5)$$

where n' are the indices of the channels selected but not completed. $C'(n)$ is assumed to be constant, it is the cost spent on any uncharacterized channel to find channel n .

The lower bound for the number of samples required to find this channel can be computed as

$$M_n + \lfloor \frac{1}{p'_L} \rfloor, \quad (5.6)$$

where we have taken the best case scenario of two samples per uncharacterized channel. We assume that the decision of whether or not to stay with a channel is made early in the sample sequence of that channel. There are only a limited set of possible paths that will yield a jump and the difference in cost between the alternative paths does not amount to more than two samples worth of cost. The critical number that governs how much cost is paid to discover a channel is the number channels searched since the M_n and $C(n)$ for the channels we did classify will be very close.

5.4.4 Choices of λ that bring the search closest to optimal

The choices of λ controls how we are penalized for expending measurements, and thus influences the utility of the greedy approach. The $F(m)$ or ζ term prevents the test from being too aggressive early on, and overlooking channels which are not far enough above or below the thresholds (p'_R, p'_L) . The quantity λ , for a given choice of test parameters

(δ, p'_L, p'_R) , is a measure of how aggressively we wish to explore for channels that are far away from the thresholds. As λ increases we spend more samples searching and fewer classifying because disagreements between samples are heavily penalized even for sequences that are very close to completion. In the ideal case, we would choose λ so that the number of samples spent searching is less than samples spent categorizing to achieve the best tradeoff between searching for channels to classify and actual classification of those channels.

When the value of λ nears the optimal value, channels that have early disagreement do not necessarily get ignored immediately. At each step the cost will force one to either pick a channel from the unmeasured pool, or continue with the channel that just received a sample. Once a channel has passed the early jump point (loosely defined as the period before the forgiveness window expires), the probability that we observe a sequence that raises to the cost high enough to cause a jump is low. This probability is low because observing a sequence that allows us to stay on a channel past the early jump point even though $p_n \in \omega_r$ is a rare event for that n . While late jumps (jumps that happen after the forgiveness window expires) can happen, they are exceedingly rare because the early jumps usually cause the system to leave most of channels that would vary that much.

With an optimal choice of parameters and the assumption of $|N| > L$ there will always be an unmeasured channel to jump to. Thus any channel that goes up in cost will never be revisited and has made as much contribution to the overall search cost as it is ever going to. Since any channel that consistently produces cost reduction will rarely turn around, the search cost that results from this parameter choice will be the lowest cost that the greedy approach can achieve. This parameter choice returns us to the case where the greedy-choice property holds.

Chapter 6

Simulation and Analysis

To verify our assumptions we ran numerical simulations over many 1000s of trials. We extracted some quantities of interest to serve as metrics for performance under various assumptions. No one choice is completely optimal nor will they work under all scenarios. We identify what the tradeoffs are for each choice we make so that actual systems built using this technique can be tuned to match the assumptions.

6.1 Validating our implementation of the SPRT

To verify our implementation of SPRT, we pick a $p \in [0, 1]$. For this p we draw a large ordered sequence of samples χ_m . Treating this sequence as the samples from an unknown parameter we compute $d_{M_\tau} = \sum_m^{M_\tau} \chi_m$ and where M_τ is the index at which this sum crosses a threshold line and causes the overall two-test approach to make a decision for trial τ . We perform this process over the entire range of allowable p (with a granularity of 0.001), and for each p we run 1000 trials ($\tau \in [0, 1000]$). For each trial with parameter p we compute an indicator function

$$I_p(\tau) = \begin{cases} 1 & \text{this trial declared } H_0 \\ 0 & \text{otherwise} \end{cases} \quad (6.1)$$

We then compute the sample mean of this indicator over all the trials for each p . If our test implementation was correct then $L(p) \approx \frac{\sum_\tau I_p(\tau)}{\max(\tau)}$ and $E_p(M) \approx \frac{\sum_\tau M_\tau}{\max(\tau)}$. In figure (4.1) and figure (4.2), this is how the simulated curves were generated.

6.2 Choosing δ for the SPRT

As we saw in section (4.11) for a fixed p'_L , the choice of δ is a trade off between effective error rates $(\alpha_{effective}, \beta_{effective})$ and $E_p(M)$. From figure (4.14) we see that there are diminishing returns for increasing deltas. For each delta as, $p \rightarrow 0$ or $p \rightarrow 1$, the $E_p(M)$ gets closer to the edge. In the range of $0.10 \leq \delta \leq 0.15$ the $E_p(M)$ clusters very tightly at the edges, but the edges only tell half the story. Because of our choice of δ and (p'_L, p'_R) , $\max_p(E_p(M)) = E_{p'_L}(M)$. While the edges are clustering, the maximum is decreasing (see figure 4.14). This maximum plays a role in the effectiveness of our search because it bounds the amount of samples we can waste on a channel.

To see the effect of δ on the search consider figure(6.1). We see that as δ increases the threshold lines themselves get closer together. This is how large δ yields shorter tests, the random walk does not have to travel as far to make a decision. The further apart the threshold lines are the lower the probability that a random walk will “accidentally” cross them. This is why larger δ makes more mistakes, it is easier to walk through the lines by chance. The narrowing of the threshold lines changes the behavior of our cost functions values, and the choice of λ that will bring us back to the optimal greedy regime.




Figure 6.1: Single Test threshold lines for varying delta.

Unfortunately, we end up with a dilemma where the choice of δ depends on λ , and

the choice of λ depends on δ . To resolve this we fix $(\alpha_{effective}, \beta_{effective})$ and determine which δ s achieve a reasonable compromise between $\max_p(E_p(M))$ and these constraints on the error. In figure (6.2) and figure (6.3) we see that a δ between $0.1 \leq \delta \leq 0.125$ keeps the $\beta_{effective} \leq 0.05$ while keeping the $\max_p(E_p(M)) \leq 200$. Given this chosen bound on the effective error rates for one decision, we can then choose a λ that pairs with this delta range.

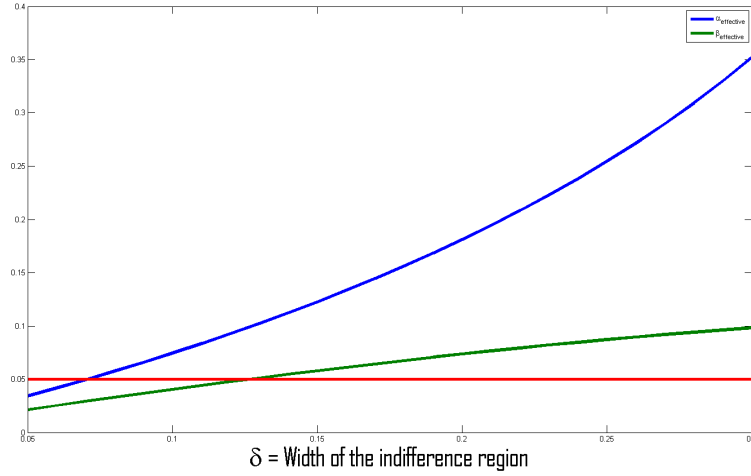


Figure 6.2: Effective error rates for varying delta.

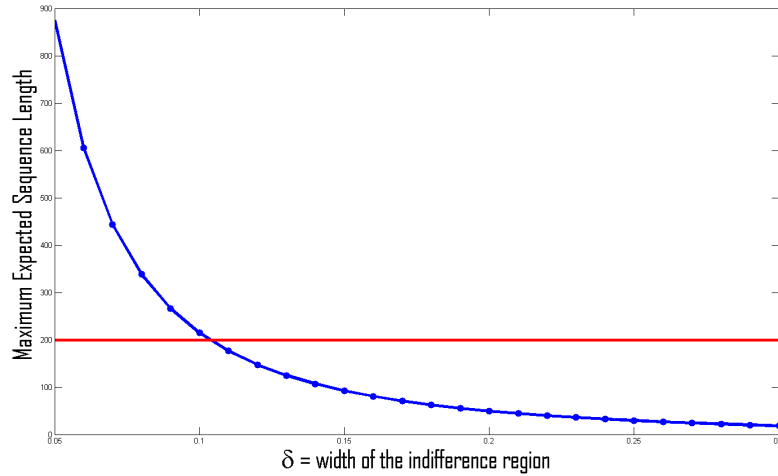


Figure 6.3: $\max_p(E_p(M))$ for varying delta.

6.3 Effects of ζ

Once we have chosen a δ , we can translate the indifference region to minimize $E_{p_n}(M)$ and make the threshold lines more fair. This is done by choosing an appropriate ζ . This translation essentially trades $\mu(\omega_a)$ for $\mu(\omega_r)$, which will ultimately impact our $\alpha_{effective}$ and $\beta_{effective}$. In figure (6.4) we see the effect of the translations on the operator curve. The effects of these translations on $E_{p_n}(M)$ can be seen in figure (6.5). In figure (6.6) we drew the threshold lines from the two-test case and the ideal




Figure 6.4: $L(p)$ for varying ζ with $\delta = 0.125$.




Figure 6.5: $E_{p_n}(M)$ for varying ζ with $\delta = 0.125$.

trajectories which result in a decision of H_1 or H_0 . We can see that the translations result in changes to the threshold line's slopes. For decreasing ζ we see a decrease in the slopes, but we also observe a decrease in the distance of separation of the threshold lines. This is true even though δ is held constant. These two changes explain why $E_{p_n}(M)$ is getting smaller. When the slope is closer to zero, there is less bias for one hypothesis vs another. Since we chose the parameters to preserve symmetry, lower the slope means the other threshold lines get closer together. This makes the two-test case more forgiving of early mistakes. As ζ decreases the $E_{p_n}(M)$ is moving towards the edges and $\max(E_{p_n}(M))$ is getting smaller. Since the expected sequence length over the decision region we are interested in is given by

$$E(M) = \int_{\omega} E_{p_n}(M) dp,$$

the shift in this $E_{p_n}(M)$ will result in smaller expected sequence lengths for those regions.




Figure 6.6: Threshold lines for varying ζ with $\delta = 0.125$.

6.4 Optimal λ for the search

In practice the lower bounds on searching calculated in section (5.4.1) are actually quite loose. The idealization required to get the greedy-choice property to hold for every allocation are not easily achievable. We can however, pick the λ that achieves a

good compromise between searching and classifying by evaluating how many samples are spent to find a channel. We can then see how this generalizes to the full search by looking at the fraction of channels classified and the error rates of those that are classified.

6.4.1 λ effect on searching for channels

At each tick of the global clock the search tries to decide if the channel it just sampled has a short $E_{p_n}(M_n)$. The cost function from section (5.1) tries to predict this by considering the distance needed to finish and the samples already consumed. Our H'_0 two-test hypothesis from Eq. (4.54) is that the parameter p_n is in one of the outer regions of the range $[0, p'_L) \cup (p'_L, p'_R) \cup (p'_R, 1]$. From what we have seen from our plots of $E_{p_n}(M_n)$ this assertion goes hand in hand with the assertion that the $E_{p_n}(M_n)$ is low. Samples spent on channels for which H'_1 are true are not of high utility because these channels have the “undesirable” characteristics that their probability of being occupied is too close to uninformed. We would like to spend fewer samples on these channels if possible. The bias in our cost function tries to take advantage of the relationship between the hypothesis and $E_{p_n}(M_n)$. Since our goal is to characterize as many channels as we can while maintaining a bound on the error, our best course of action is to focus most of our samples on channels where H'_0 is true.

The search procedure will hop around, measuring unexplored channels, in an attempt to find channels where $E_{p_n}(M_n)$ seems short. Once it deems one promising, it then places enough samples in this channel to get a completed test. We control what is deemed promising by adjusting the cost of measurements, λ . If measurements are expensive, high λ , our search will hop frequently and very often overlook channels that may actually satisfy H'_0 as noted in section (5.4.1). This has two implications, the first is that we will spend many more samples searching. The second is that in order to generate enough cost reduction to offset the measurement cost increase, the p_n that does capture the interest of the search will be very far from the thresholds. On the contrary for low λ we are more tolerant to mistakes and will accept values of p_n closer to the thresholds. Thus we will spend less samples searching, however, we will spend

more samples characterizing.

To examine this trade off, we consider how many samples are required to find one channel for which H'_0 is true. With the δ from section (6.2), we can fix all other parameters at $\delta = 0.125, p'_L = 0.2, p'_R = 0.8, N = 1024$ and allow λ to vary. From section (5.2), our λ is measured in fractions of a sample $\frac{1}{2}$, and thus is only meaningful in the range $[0, 1]$. For each lambda we run the search 1000 times and collect some meaningful quantities.

In figure (6.7), we see the number of samples required to categorize the channel we found as well the number of samples spent searching for that channel as a function of λ . With the given parameter choices, the minimum number of samples to characterize a channel is 25. As λ increases the number of samples required to characterize reaches this limit but the number of samples required to find the channel with this property grows disproportionately large.




Figure 6.7: Samples required to find one low $E_{p_n}(M_n)$ channel and Samples required to categorize that channel vs λ

To confirm our conclusion, in figure (6.8) we plotted the mean across trials of the actual parameter value found as a function of λ for the left side. It decreases with λ as expected. This drop in the p_n is not very large, and for individual channels it is not enough to justify the large extra cost. We can see a fundamental trade off between spending samples to explore and spending samples to categorize. It is worth noting

that even with $\lambda = 0$ the p_n is very far from p'_L .




Figure 6.8: Actual value of the paramter p_n found vs λ

6.4.2 How often do I pass on channels I should have stayed on?

To consider the efficiency of our search procedure, we would like to examine the amount of samples that are “wasted”. From figure (6.7) we have an idea of how many samples went into finding an individual channel, however this number hides some of the details of how we get to that decision. In examining the amount of samples spent we can partition the set of channels we passed through to get to the decision in two sets, the channels we should have accepted ω'_a and the channels we should have passed on ω_p . The passed channels samples are actually not “wasted” in some sense as they are the cost of searching. The should have accepted set however does represent wasted samples as we should properly have settled on these channels.

In an ideal case $|\omega'_a|$ should be small. The obvious question is, how does a channel end up in this set? As we saw from figure (4.2), when the parameter are examining is inside the indifference region $\omega_I = \{p|p_0 \leq p \leq p_1\}$, expected sequence length grows very long. Channels with parameters in this region will produce random walks that drive up costs above those of an unmeasured channel. When our search lands on one of these channels it is very likely that the search will jump to a different channel because

samples of this channel produced a sequence that has a cost that is higher than the unexplored cost.

Suppose we specify the resource regime $\frac{L}{|N|}$, the assumed meta distribution, and the threshold choices (p'_L, p'_R) . If we take as a criteria of admissibility, the maximum probability of error $(\beta_{effective}, \alpha_{effective})$, the parameters of λ , δ and ζ can be chosen to lower $|\omega'_a|$. This quantity is an important measurement tool for determining how many samples are spent exploring and how many are spent categorizing (exploiting).

While we do not want this quantity to be large, the case of $|\omega'_a| = 0$ is also not efficient. If $|\omega'_a| = 0$, we would never pass on channels near the threshold, and would thus spend $\max(E_{p_n}(M_n))$ samples to categorize some of them. Suppose the search landed on channel n with parameter very close to p'_L or p'_R but properly on the H_0 side of the threshold, e.g. $p_n = p'_L - \epsilon$. $\max(E_{p_n}(M_n))$ is orders of magnitude higher than channels with parameters below the indifference boundary, e.g., $n' \neq n$ with $p_{n'} < p'_L - \delta$. As long as N is large, we may actually characterize more channels if we continue to search instead of trying to characterize channel n even though n properly below the threshold. Our choice of λ, δ and ζ needs to strike a balance between $|\omega'_a|$ and $|\omega_a|$, the set that we failed to classify but were below the threshold, and the set we did classify. In figure (6.9) we can see the quantity of channels passed on for varying δ and multiple choices of λ . We will want to choose low λ and δ in the previously specified range to prevent $|\omega'_a|$ from growing too big and consuming most of our samples searching instead of classifying.

6.5 Alternative schemes for comparison

We will examine some very simple schemes that can be computed very fast. Each of these schemes has a draw back over the greedy-SPRT, however there are resource regimes, where the increase in probability of error is small, and the gains in computational speed greatly out weigh the greedy-SPRT approach.

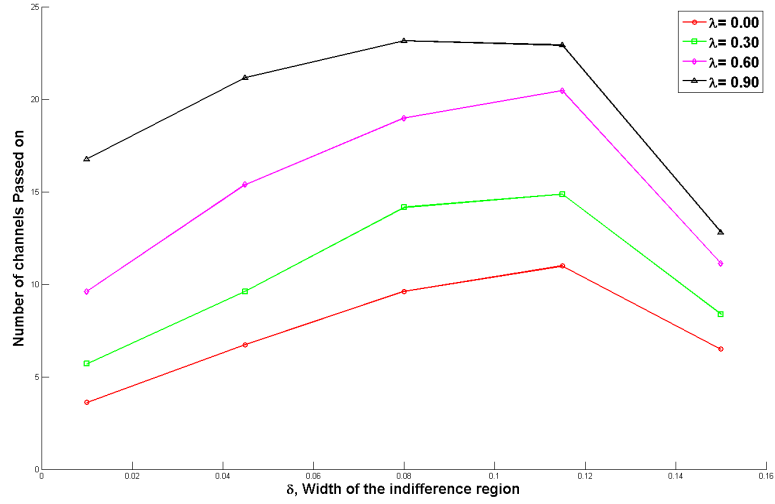


Figure 6.9: Channels passed on as a function of λ and δ

6.5.1 The simplest allocation scheme

The most simple allocation scheme one can consider distributes $M_n = \frac{|N|}{L}$ samples to each channel (in no specific order). It repeats this process until it runs out of samples. It then computes the estimate of the parameter $\hat{p}_n = \frac{\sum^{M_n} x_{n,m}}{M_n}$ based on the samples it has distributed to channel n . Using each estimate, it then decides which hypothesis to claim by comparing \hat{p}_n to the thresholds.

Clearly M_n is fixed for every channel and more importantly increasing $N \implies$ decreasing M_n . This has the apparent problem that as M_n decreases both $(\alpha_{effective}, \beta_{effective})$ increase. In section (6.6.1), we will compute the error rates for the declarations we have made.

There are two choices of how to distribute the samples for this method. We can scan from lowest index to highest index, and then wrap around until we run out of samples. The alternative is to start from the lowest index, place M_n samples in this channel, and keep moving along until we run out of samples. There may be practical considerations for which technique we would want to follow, e.g. the second method makes the estimates available for immediate use, but might introduce a time correlated bias. Under our assumptions of fixed p_n for the measurement duration, however, both techniques produce equivalent decisions, and are thus identical. We will compare this

scheme to the greedy-SPRT scheme in section (6.7).

6.5.2 A tree structure approximation

In the greedy-SPRT case we are using the distance to threshold lines at each step as a guide for allocation of measurements to channels. Each individual sample is allocated to minimize the time spent searching and maximize the number of channels characterized. In the regime of small $\frac{L}{N} < 1$ we will never be able to cover all possible channels, and so by prioritizing channels with the shortest expected sequence length $E_{p_n}(M_n)$, we try to achieve a compromise that yields the desired maximization.

We can modify the simple strategy, to be adaptive by building tree with a process that alternates between breadth and depth searches. We will build the tree in phases. At each phase we employ a pruning rule to prevent us from placing too many samples into a channel that confirms the wrong hypothesis, or takes too long to confirm the right hypothesis. This rule will be based on the agreement (or lack thereof) of the samples.

First we will need to decide how many samples we place in a single channel per phase, name this quantity ψ . The smallest amount we can choose is $\psi = 2$, because $\psi = 1$ samples can not be checked for agreement and parameter estimate $\hat{p}_n \in \{0, 1\}$ does not really give us anything to decide on. The choice of ψ will be discussed in section (6.5.3).

Next we will set aside a fraction ρ of our sample budget to preform the initial pass. We allocate ψ measurements per channel until we exhaust our first pass budget ρL . Since the channels are not in any order, simply going sequentially by index selects a random sampling. This first pass will place measurements in $|N_1| = \frac{\rho L}{\psi}$ channels. After this first pass, we inspect the elements of N_1 for agreement, and establish a criteria for admissibility that picks the set we will sample in the second phase.

To build the admissibility rule, consider the case of $\psi = 4$. We can enumerate all 2^4 choices for the samples of an individual channel n , but we know that $\hat{p}_{n,1} = \frac{\sum^{M_{n,1}} x_{n,m}}{M_{n,1}}$ is a sufficient statistic. That is we do not need to consider the ordering of the sequence of 1s and 0s that results from sampling, merely the sum of it's elements. We have given each of the parameters an additional subscript to indicate what phase they are from, e.g.

$M_{n,1}$ is the stopping time for channel n in phase one. In this case, $\hat{p}_{n,1} \in \{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\}$ corresponding to the cases of 0, 1, 2, 3, 4 ones.

We can pick a rule for agreement of samples which will translate to a threshold on $\hat{p}_{n,1}$. As an example, we will allow for one disagreement among the samples of an individual channel. That is at least three zeros or three ones in the set of four samples for this example. This will admit all channels where $\hat{p}_{n,1} \neq \frac{1}{2}$, call this set of channels admitted during the first pass $N_{1,a}$. For each channel in $N_{1,a}$ we repeat the process distributing ψ samples to each channel in $N_{1,a}$. At this phase every channel in N_2 has eight samples. Now $\hat{p}_{n,2} \in \{0, \frac{1}{8}, \frac{1}{4}, \frac{3}{8}, \frac{1}{2}, \frac{5}{8}, \frac{3}{4}, \frac{7}{8}, 1\}$, we can use the same admissibility rule which would limit us to two errors in eight samples. Thus all channels with $\frac{3}{8} \leq \hat{p}_{n,2} \leq \frac{5}{8}$ would be excluded. Our admissibility rule of one in four disagreements essentially admits channels with $\hat{p}_n \leq 0.25$ or $0.75 \geq \hat{p}_n$ to the next phase.

Proceeding to phase three we apply the same process onto the set $N_{2,a}$, and continue the same way with each phase until we exhaust our budget. We can then categorize all channels that have received samples by simply computing \hat{p}_n for each channel and then comparing that to the given thresholds (p'_L, p'_R) . There is one early stopping condition that might cause this process to break, for some phase ν , we may have that $N_{\nu,a} = \emptyset$. There are a few ways to resolve this issue, one can relax the requirements on disagreement for what ever phase you are at, e.g. allow three errors in eight at phase two instead of two. This has the effect of increasing the error because we are relaxing our admission threshold, but also means we will not “throw away” some of the early samples we spent. Alternatives are considered in section (6.5.4). We note that if we set $\rho = 1$ and $\psi = \frac{L}{N}$ this scheme reduces to the simple scheme of section (6.5.1).

A pseudo code implementation would look like:

```
#First Pass
MaxScanIndex = (rho * L) / psi
for i = 1 to MaxScanIndex
    d(i) = d(i) + sample(i, psi)
    m(i) = m(i) + psi
end
```



```

#Spend the rest of the samples
UsedSamples = (rho * L)
While UsedSamples < L

    #reset picked counter
    ChannelsPicked = 0

    #Sample any channels that meet criteria
    for i = 1 to MaxScanIndex
        pHat(i) = d(i) / m(i)
        if {(pHat(i) < LeftThreshold)
            or (RightThreshold < pHat(i))}
            d(i) = d(i) + sample(i, psi)
            m(i) = m(i) + psi
            UsedSamples = UsedSamples + psi
            ChannelsPicked = ChannelsPicked + 1
        end
    end

    #Relax the threhsold if we didn't find any thing
    if ChannelsPicked = 0
        LeftThreshold = LeftThreshold - RelaxationAmmount
        RightThreshold = RightThreshold + RelaxationAmmount
    end

end

#Now compute the estimates and
place them in their proper characterized sets.

```

```

for i = 1 to MaxScanIndex
    pHat(i) = d(i) / m(i)
    if pHat(i) < LeftThreshold
        add i to set N_Left
    else pHat(i) > RightThreshold
        add i to set N_Right
    else
        add i to N_Center
    end
end
end

```

Where the N_{Left} , N_{Right} and N_{Center} are sets that contain channel indices that satisfy specified conditions.

6.5.3 Parameter Choices for the tree approximation

This scheme is controlled by the two parameters (ψ, ρ) and the admissibility rule. For a fixed ρ as ψ increases we might expect to get lower overall error rates because all the poorly explored channels are better approximated, but this increase in ψ comes at the cost of smaller $|N_1|$. While we may reduce the probability that we place a channel n into the wrong region (ω_a, ω_r) , we also shrink the size of channels we can say anything about. This may not seem like a limitation however the fraction of channels that will confirm H_0 for the set N_1 has cardinality $2|N_1|p'_L$ in the symmetric hypothesis case. While we do not have an $E_p(M)$ curve for this scheme, the Bernoulli parameter observation of section(4.13) still holds. As $|N_1|$ decreases, we will find few of the desirable channels, and thus fewer have successful categorizations. As before we need to choose (ψ, ρ) to achieve a balance between samples spent exploring, and samples spent exploiting.

We can also set aside a larger ρ at the beginning, this will cast a wider net initially. This results in a higher $|N_1|$, which yields higher probabilities for landing on a channel that will confirm H_0 . The downside to this is that we will have fewer reserve samples to use for later phases. Thus our trees cannot grow very deep, and the number of

measurements spent classifying channels of interest is lower. As we proceed to later phases the granularity of values of the parameters estimated from the longest branches of the tree, $\hat{p}_{n,final}$, can take on in the final phase increases. The more phases we go through, the more discerning our test of the parameter will be for channels that make it to the final phase. Channels with parameters very close to the threshold may get misclassified if the granularity of $\hat{p}_{n,final}$ is not fine enough to make a proper decision. This can happen if we run out of samples to achieve the necessary granularity because we spent too many in the first phase.

The admissibility rule is closely tied to our choice of (ψ, ρ) . We can choose an arbitrary admissibility rule, but if the admissibility rule and (ψ, ρ) do not agree the test fails to meaning fully distribute samples in an adaptive way, and basically reduces to the simple scheme. As an example we could have considered an alternative ψ choice of $\psi = 5$, then our admissibility rule could be one in five errors, which would match with our choice of (p'_L, p'_R) . This would be a perfectly valid choice, but may be too aggressive for some resource regimes.

6.5.4 Alternative tree approaches

In building the tree implementation, we made a few design choices that have some alternative options. Each of these options represents some trade off between samples spent searching samples and spent classifying. We did not opt to implement these designs as they would have raised the complexity of the tree building process. It may be possible to achieve tighter approximations to the greedy-SPRT case but at the cost of increased scheme complexity.

The first alternative to be considered is how to handle the $N_{\nu,a} = \emptyset$ case. Our simple approach was to relax the criteria on measurement disagreement. As an alternative, we could start the problem over again with $L' = L - \sum_{i=1}^{\nu} |N_i| \psi$ samples (what ever samples were left) and pick channels from N_1^c (channels we have never measured yet). In this case, we would not relax the admissibility rule, instead we widen the search over larger sets. The gain from this approach is wider breadth of search, however since we are scanning more channels, we are also adding more misclassified channels to our set

of measured channels which raises our error.

As a second alternative, we can attempt to patch the low sampled channels problem, by setting some threshold of measurement. The smallest threshold we can pick is to discard all channels that never made it past phase 1. The major draw back to this is that all the samples of this phase will essentially be discarded. This is ultimately governed by the choice of (ρ, ψ) . We can try to reduce this effect by lowering the two parameters, but that results in a loss of parameter granularity (and thus an increase in misclassification errors).

We can also try to be more conservative with our samples by not requiring equal sample distribution across all phases. Perhaps at phase 1 we use ψ as our initial measurement allocation, then for each subsequent phase we only allocate ψ' samples where $\psi' < \psi$. The problem with this alternative is that we will require a new admissibility rule for subsequent phases. One can imagine a sequence of ψ^t , one for each phase. Then at each phase we will have a different admissibility policy to go with the ψ^t .

If we examine the element-wise decisions that are made by the greedy-SPRT approach, these alternatives taken together with per-phase ψ^t preform the similar operations to the greedy-SPRT. The greedy-SPRT randomly chooses a channel initially, and then begins allocating samples to it. If the channel shows promise by reducing the distance to the threshold lines, it is rewarded with more samples. If not, the test moves on. While the greedy-SPRT is not preformed in phases, the allocation patterns that arise from either strategy will look the same, perhaps with a reordering or initial choices.

6.6 Performance Metrics

In general a global performance metric is hard to derive because there are many tuning parameters and many scenarios where the choice of parameters leads to an unworkable solution. For example, we could spend all the samples searching and make no declaration, or we could blindly allocate all samples to one channel, yielding the lowest error but only one decided channel.

For any sample distribution scheme to be admissible it would have to maintain bounds on the effective error rates. The ideal metric for an admissible scheme would be number of channels characterized while maintaining the bound. However the error bound itself is not clearly defined. This metric is itself hard to compare in closed form because for some of the evaluated schemes this quantity is a function all the parameter choices and adjustments we make to the scheme to make it admissible.

We are going to examine a somewhat flawed but useful error metric which we can compute very easily. The performance metrics we can examine will be samples spent finding a H_0 confirming channel, channels passed on during the search and the number of channels found at the end of the search. All of these metrics are directly computable from simulation, but may not necessarily admit a closed form expression which is derivable from the model.

6.6.1 Error Probability Across the entire search

For the error of any individual test, we will take a simple approach to quantify the errors made across the search. If we consider the set N_{H_0} of channels for which we have declared H_0 , any particular element of this set $n \in N_{H_0}$ is incorrect with probability $\beta_{effective}$ by definition. Since each element of N_{H_0} is placed in there independently of the others, we can treat them as an independent population with the property that some of them may be wrongly categorized. Let N'_{H_0} be the set of channels that are wrongly categorized. To determine the probability that any element of this population is wrongly categorized we simply compute $P(n \in N'_{H_0}) = \frac{\mu(N'_{H_0})}{\mu(N_{H_0})} = \frac{|N'_{H_0}|}{|N_{H_0}|}$, where μ is the counting measure. Let $\beta_{search} = \frac{|N'_{H_0}|}{|N_{H_0}|}$. In a similar way we can compute an expression for α_{search} .

While this derivation is perfectly reasonable, it does gloss over some important facts. These error rates given here only tell you about the channels we did classify. We do not count all the samples lost searching. In the SPRT case we can compute $\beta_{effective}$ directly from the test itself via the integrals in section (4.11). There is a problem, however, with directly equating these two quantities. The act of searching terminates any tests that takes too long, and as we see from the $E_P(M)$ curve of figure (4.11) and

the OC curve of figure (4.12), these are the tests that have higher probability of error. Consider the integral of Eq. (4.41) and note that for our test the search prevents us from picking parameters in the $[p', p_1)$ region. Thus we can bound our effective error by $\beta\mu(\omega_r) = \beta\mu([p'_L + \delta, p'_R - \delta]) = \beta(p'_R - p'_L - 2\delta)$, where we have used the fact that the measure of the interval is the difference of its end points (Lebesgue). Here ω_r is the region between the two thresholds because of how our hypotheses are chosen. We expect that $\beta > \beta(p'_R - p'_L - 2\delta) \geq \frac{|N'_{H_0}|}{|N_{H_0}|} = \beta_{search}$ which is observed in the simulations of section (6.7).

For the other schemes we do not have bounds on the effective error. As we shall see in section (6.7) $\frac{L}{|N|} \searrow \implies \beta_{search} \nearrow$ for both alternative schemes. The error rates of channels that were classified fail to consider the efficiency of the search. In many cases we might erroneously leave a channel when the parameters might have been close to what we were looking for. These types of mistakes do not show up in these error calculations because these channels never reach a decision. Our goal of maximal number of channels classified can be achieved by bounding the errors described here, while maximizing the metric described in section (6.6.2). This maximization will be achieved by minimizing the “unreported” errors caused by passing on channels we should have stayed on, like those considered in section (6.4.2).

6.6.2 Number of Channels classified

The number of channels classified is really the ideal metric of performance, but by itself it is not a complete picture of the performance of the search. Consider the simple scheme when $\frac{L}{|N|} = 1$. At this resource regime the simple scheme manages to classify every channel but has only one sample per channel. The simple scheme will make many errors because with only one sample $\hat{p}_n = \frac{\sum_{m=1}^{M_n} x_{n,m}}{M_n} = x_{n,1}$ and thus $\hat{p}_n \in \{0, 1\}$. When we threshold on these estimates, our decisions have a very high probability of error due to the low sample rate. Since we only take one sample, when the test declares H_0 the probability that it is wrong can be computed from $P(H_0|H_1) = P(x_1 = 0|p \in \omega_r) = E_p(\phi(X)) = \int_{\omega_r} dP_p(x) = \mu(\omega_r)$ due to the definition of $E_p(\phi(X))$ given in Eq. (4.4) and that $\phi(x_1 = 0) = 1$.

We can compute an idealized upper bound on what can be discovered by any adaptive scheme by treating the search as a geometric random variable as we did in section (5.3). Since the parameters are drawn from a uniform distribution, the success probability is merely $\mu(\omega_a)$ which in the case of symmetric thresholds is $2p'_L$. The expected number of channels before we find one is the $\frac{1}{2p'_L}$. If we assume that each undesirable channel only consumes 2 samples (a perfect jump every time), then we will expend $\frac{2}{2p'_L} = \frac{1}{p'_L}$ samples to find a good channel. Once this good channel is found it will take M_n samples to characterize. Let us assume that $\forall n, M_n = M_{min}$, then $\max(|N_{char}|) \sim \frac{L}{M_{min} + \frac{1}{p'_L}}$. For our simulation scenario the SPRT parameters we have chosen yield $M_{min} = 19$ and the bound becomes $\frac{1000}{19 + \frac{1}{.2}} \approx 41$.

An important characteristic of this bound is that it does not depend on N at all. Consider N' with $N' > N$, but with L fixed. The resource regime is now worse but our strategy remains the same. The only reason we need the assumption of large $|N|$ is so that there will always be an channel with $C(0, n) = M_{min}$, the unmeasured cost. The unmeasured cost is always starts off as the minimum samples to characterize. As a channel receives samples it must then require fewer samples to characterize than an unmeasured channel in order to continue receiving samples. Assuming N' is drawn from the same distribution as N , larger $|N|$ does not change our strategy. As we will see in section 6.7, the set of discovered channels levels off after N of a certain size.

For the simple scheme, it will always characterize a deterministic number of channels given by

$$|N_{char, simple}| = \begin{cases} |N| & \text{if } |N| < L \\ L & \text{if } L < |N| \end{cases}$$

Depending on the implementation of the tree scheme the amount of channels is characterized is variable. Using the implementation we initially described, however, the number of channels characterized $|N_{char, tree}| = |N_1| = \frac{\rho L}{\psi}$. Clearly both schemes will characterize as many channels if not more than greedy-SPRT, but to achieve this result they sacrifice error probability, often beyond tolerable limits.

For the greedy-SPRT approach we can examine the how the samples are spent to get an idea of where our channel classification will fall with respect to the computed bound.

As the number of samples required to find a channel with low $E_{pn}(M_n)$ quantities grow, the $|N_{char,greedy-SPRT}|$ decreases, through our choices of the parameters (δ, λ) , we achieve the best compromise between the allowed error rates and the discovered channels.

6.7 Comparison of performance between different allocation schemes

In figure (6.10) we can see the number of channels discovered for all 3 schemes as a function of the number of channels we have to search, N . The simple scheme excels at discovering channels, however as we see in figure (6.11) this comes at the cost of high classification error.




Figure 6.10: Number of channels found for all 3 schemes as functions of N

Because the simple scheme involves large numbers, the graph scales hide the details of how the tree scheme compares with the SPRT. The tree scheme is an adjustment to the simple scheme that does not add much computational complexity, at least compared to the greedy-SPRT scheme. If we look at figures (6.12) and (6.13) we can see how the tree and greedy-SPRT compare. Clearly the tree does a better job of finding more channels, it almost approaches the bound we derived in section (6.6.2). However figure (6.13) shows us the cost of this wider exploration, higher error. We originally picked our $\beta = 0.05$, this bound is met by the greedy-SPRT scheme.

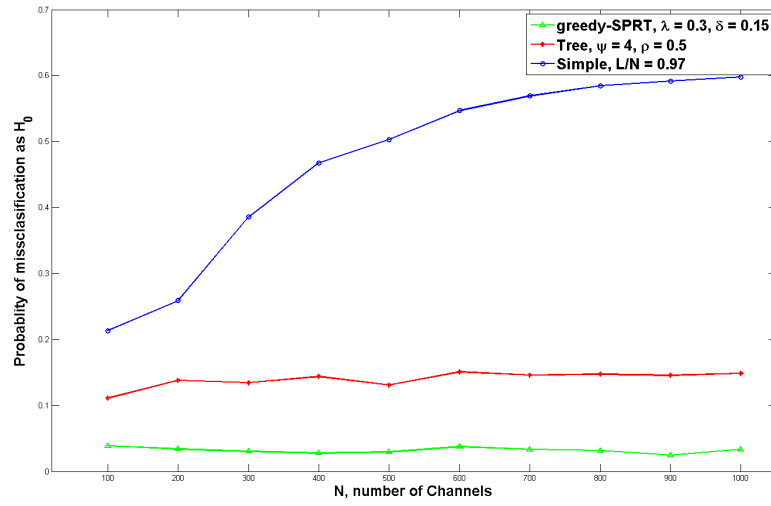


Figure 6.11: β_{search} for all 3 schemes as functions of N

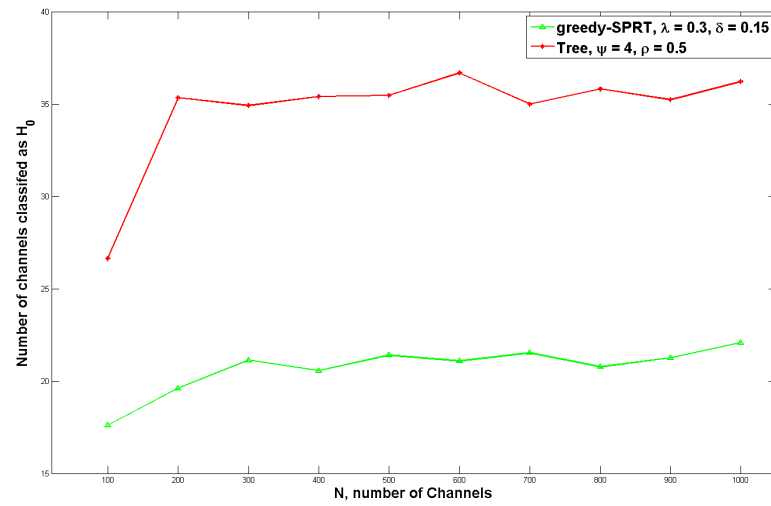


Figure 6.12: Number of channels found for greedy-SPRT vs Tree schemes as functions of N




Figure 6.13: β_{search} for greedy-SPRT vs Tree schemes as functions of N

6.8 Measurement Distribution

We can check the distribution efficiency of the SPRT as a function of N . Figure (6.14) gives us a count of how many channels we passed on. As N grows, this quantity increases as we have more channels to jump to. There is a saturation point near 400 channels after which the variation begins to dampen. After this point the size of N begins to matter less as the constraint due to L becomes a bigger factor in the curves behavior. Figure (6.15) shows some interesting results as N grows which helps explain the behavior of figure (6.14). We see that the number of samples to complete the test reaches an inflection near 400. This inflection is at the minimum number of samples required to complete a test. This is expected as the number of channels which are far below the threshold grows with N due to the uniformity assumption on the distribution of p . While there is some jitter in the number of samples to find a channel with low $E_{P_n}(M_n)$, the swing is only ± 3 samples. Here we see that for large N we spend almost as many samples searching as we do classifying.

To get a better idea of how the schemes compare with respect to the sample distribution consider figure(6.16). This histogram shows the mean number of channels that received a given amount of samples for the low $\frac{L}{N}$ regime. The simple scheme blindly




Figure 6.14: Number of channels passed on as a function of N

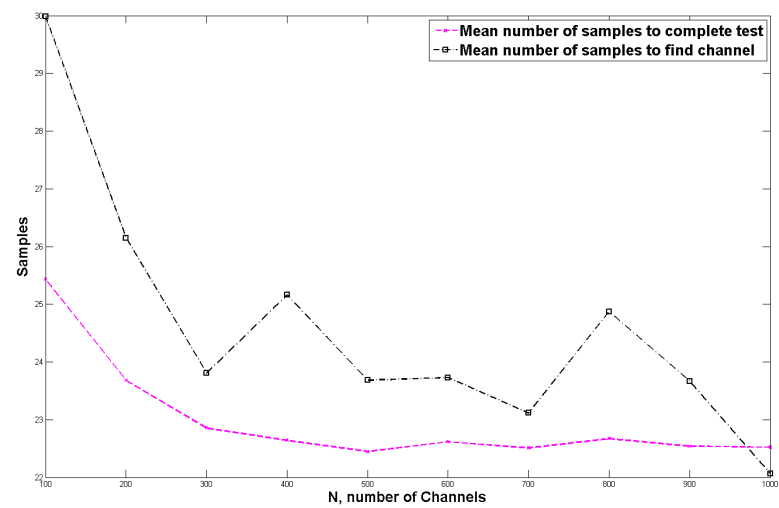


Figure 6.15: Mean samples spent classify vs searching as functions of N

puts 1 sample in almost every channel because none of it's channels revived 0 samples and the number of channels that received 1 is off the chart. The Tree Scheme is more judicious in it's distribution, and ignores a large section of channels as the 0 is off the chart. Here we see how the tree approximates the greedy-SPRT which has a decreasing but wide range of distributions. The huge jump at 19 is the result of successful searching, as 19 is the minimum number of samples for this choice of parameters. There is also a large unmeasured set, and spots of samples spent on channels slightly more difficult to characterize, e.g. 23 and 27.

We also see that the greedy-SPRT is eliminating channels with as few as 2 samples, while the tree always requires a minimum of 4. When the tree completes it's first phase it spends $\approx 75 * 4 = 300$ samples to eliminate that set of 75 channels, while the greedy-SPRT eliminates $\approx 76 = 40 + 21 + 15$ channels in $\approx 40(2) + 21(3) + 15(4) = 203$ samples. The greedy-SPRT is more efficient at finding channels worth examining as it uses fewer samples to eliminate the same amount of bad candidates.




Figure 6.16: Mean number of channels that receive a given level of sampling

Chapter 7

Conclusion

7.1 Why a proof of optimality is hard

To properly prove the optimality of the greedy-SPRT scheme one would need to derive a meaningful metric that was free of ambiguity and special cases. In the way this problem was formulated the metric of choice seems to be maximizing $|N_{char}|$ while keeping $(\beta_{search}, \alpha_{search})$ minimized. A few problems that arise with this formulation. The first is that a proper optimization of the SPRT requires optimization of a vector of parameters $(\alpha_R, \beta_R, p_{1,R}, p_{0,R}, \alpha_L, \beta_L, p_{1,L}, p_{0,L}, \lambda)$ in the greedy-SPRT case to find a balance between β_{search} and $|N_{char}|$. We have tried to simplify the parameter space by appealing to symmetry but we loose some flexibility in the class of tests we can consider when we do this.

Second, this metric is heavily tied to the distribution assumed for p_n . The sensitivity of the metric to this distribution might give rise to special cases that would prevent any scheme from reaching global optimality. This can possibly be addressed by confining the slices of spectrum we are interested in to sufficiently small pieces so that the uniformity assumption holds, but if we do that we may end up in the high $\frac{L}{N}$ regime, in which case the choice of scheme might be moot.

Third, In section (6.6.2) we used idealized assumptions to show what a bound on the channels discovered would be while keeping error rate fixed. To properly establish this bound we would need to specify M_{min} via some external criteria. Once we knew this quantity we would then need to derive expressions for the number of channels found for this scheme and show that it gets closer to this bound over any other scheme. These expressions are not necessarily trivial to derive because of the complex set of relationships between $L(P)$ and $E_P(M)$. Our best approach in this case would be to

appeal to dynamic programming techniques to traverse the problem backwards and see if the allocation scheme we proposed comes close to the perfect scheme.

7.2 Possible Applications

Our formulation of errors differs quite a bit from the usual interpretation of type I (false Alarm) and type II (miss). In [19] false alarms are interpreted as lost spectral opportunities, while missed detections result in collision. These interpretations are due to the simple hypothesis used in the formulation of the problem, specifically

$$H_{0,n} \quad : \quad \text{Channel } n \text{ is free}$$

$$H_{1,n} \quad : \quad \text{Channel } n \text{ is busy}$$

Since our hypothesis is formulated somewhat differently, the interpretation of our errors is also different. For our formulation type I errors are interpreted as mislabeling channels which vary a small amount as channels that alternate between occupied and not, very often. Type II errors then label relatively constant channels as volatile.

Because the system goal is different, the performance of our system has to be measured differently. One can utilize the information that comes from this system to improve the performance of the transmission system proposed in [19]. The output of the system will partition N_{char} into three sets $(N_{Left}, N_{Center}, N_{Right})$, the left, center, and right sets of characterized channels. In the worst case performance scenario, we can start a follow up system like the one in [19] and not allow it to use any channels from N_{Right} . In the most ideal case we start the follow up system and only allow it to use channels from N_{Left} .

If All the channels end up in N_{Center} we expect poor performance from the follow up system. In the absolute worst case all channels will be in N_{unchar} , under this scenario we either did not have enough samples to do anything useful, or the parameter distribution is not uniformly distributed (our assumptions are wrong). Again the follow up system will suffer and might not want to try at all (for example to conserve power). This might

be the case if the entire space of channels is populated with channels that sit on one of our two respective thresholds forcing all SPRTs to require $\max(E_{p_n}(M_n))$ samples.

7.3 Future Work

From here there are several avenues that can be explored, the first and foremost is establishing the optimality of the scheme as described in section (7.1). Additional avenues include considering combination of results across multiple scanning passes to build historical models. We can then use the historical models combined with the most recent analysis to make better informed decisions about what channels to try transmissions on.

We can also attempt to numerically evaluate a more general (possibly asymmetric) set of parameters that govern the scheme, and then pick the subset of parameters that is both admissible (yields tolerable error probabilities) and maximizes $|C_{char}|$.

Finally, all the simulations are representations of a real world radio environment. We know that uniformity of the parameter space is an approximation, but it would be useful to validate the model with real world RF. We can employ current software defined radio technology to test out our parameter assumptions on real world channels. While real hardware has it's own complexities, eventually any proposed system needs to be tested on actual spectrum.

7.4 Final Remarks

In figure (7.1) we have a graph of the mean computation time per trial in seconds vs number of channels. While this number is heavily dependent on implementation details, it does demonstrate another key distinction between the greedy-SPRT and the other two schemes. The greedy-SPRT is very computationally expensive when compared to the other two schemes. In high resource regimes, i.e. high $\frac{L}{N}$ the gain in error reduction might not warrant the expense of computational cost. As $\frac{L}{N}$ decreases the greedy-SPRT maintains a bound on the error by spending many more samples and many more computations to find channels that have shorter $E_{p_n}(M_n)$ which it can then classify

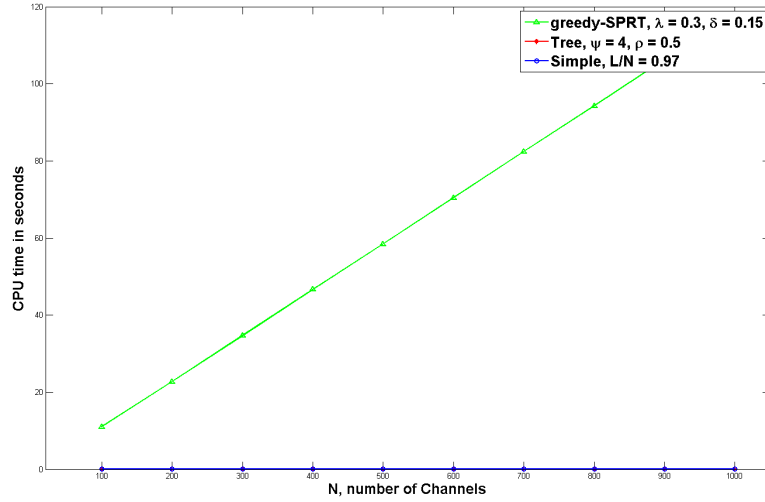


Figure 7.1: CPU Time for All 3 schemes as function of N

with a relatively short sequence.

One of the reasons this is true is that the parameters were tuned for a very low resource regime. The choice of parameters can be adjusted to be more or less judicious with the expenditure of samples. If we knew that the resources were more abundant, a different choice of parameters might actually lower the computations by reducing the amount of search required to find a channel to dwell on.

Another feature that is not captured in any of these graphs is that in the high resource regime, the greedy-SPRT often completes early, using fewer samples than it was allotted to complete the task. The other schemes are built to exhaust the budget and thus cannot adapt to this condition. Even in the high resource regime, one might prefer the greedy-SPRT if we wish to conserve measurement resources for a followup system.

References

- [1] E Biglieri; A. Goldsmith; L. Greenstein; N. Mandayam; H. Poor. Principles of Cognitive Radio. Cambridge University Press, Cambridge, UK, 2013.
- [2] Tefvik Yucek and Huseyin Arslan. A survey of spectrum sensing algorithms for cognitive radio applications. Communications Surveys & Tutorials, IEEE, 11(1):116–130, 2009.
- [3] Joseph Mitola III and Gerald Q Maguire Jr. Cognitive radio: making software radios more personal. Personal Communications, IEEE, 6(4):13–18, 1999.
- [4] Anita Garhwal and Partha Pratim Bhattacharya. A survey on dynamic spectrum access techniques for cognitive radio. arXiv preprint arXiv:1201.1964, 2012.
- [5] Hongjian Sun, Arumugam Nallanathan, Cheng-Xiang Wang, and Yunfei Chen. Wideband spectrum sensing for cognitive radio networks: a survey. Wireless Communications, IEEE, 20(2):74–81, 2013.
- [6] Ian F Akyildiz, Won-Yeol Lee, Mehmet C Vuran, and Shantidev Mohanty. Next generation/dynamic spectrum access/cognitive radio wireless networks: a survey. Computer Networks, 50(13):2127–2159, 2006.
- [7] Judith A Spray. Multiple-category classification using a sequential probability ratio test. Annual Meeting of the National Council on Measurement in Education, 93(7), 1993.
- [8] Jay Bartroff, Matthew Finkelman, and Tze Leung Lai. Modern sequential analysis and its applications to computerized adaptive testing. Psychometrika, 73(3):473–486, 2008.
- [9] Lifeng Lai, Hai Jiang, and H.V. Poor. Medium access in cognitive radio networks: A competitive multi-armed bandit framework. In Signals, Systems and Computers, 2008 42nd Asilomar Conference on, pages 98–102, 2008.
- [10] Lifeng Lai, H. El-Gamal, Hai Jiang, and H.V. Poor. Cognitive medium access: Exploration, exploitation, and competition. Mobile Computing, IEEE Transactions on, 10(2):239–253, 2011.
- [11] Yan Xin and Lifeng Lai. Fast wideband spectrum scanning for multi-channel cognitive radio systems. In Information Sciences and Systems (CISS), 2010 44th Annual Conference on, pages 1–6. IEEE, 2010.
- [12] Robert G. Gallager. Discrete Stochastic Processes. Kluwer Academic Publishers, Norwell, Massachusetts, 1996.

- [13] E. L. Lehmann. Testing Statistical Hypotheses. John Wiley & Sons, Inc, New York, 1959.
- [14] A. Wald and J. Wolfowitz. Optimum character of the sequential probability ratio test. The Annals of Mathematical Statistics, 19(3):326–339, 1948.
- [15] Robert G. Gallager. STOCHASTIC PROCESSES:Theory for Applications. Draft, Draft, 2013.
- [16] Abraham Wald. Sequential Analysis. Dover Books, New York, N.Y., 1973.
- [17] Milton Sobel and Abraham Wald. A sequential decision procedure for choosing one of three hypotheses concerning the unknown mean of a normal distribution. The Annals of Mathematical Statistics, 20(4):502–522, 1949.
- [18] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. Introduction to algorithms, third ed., 2009.
- [19] Lifeng Lai, Hesham El Gamal, Hai Jiang, and H Vincent Poor. Cognitive medium access: Exploration, exploitation, and competition. Mobile Computing, IEEE Transactions on, 10(2):239–253, 2011.