8-2007

Performance of Serial, Matched-Filter Packet Acquisition Algorithms with a Preamble-Sequence Acceptance Criterion

Javier Schloemann
Clemson University, sjavier@clemson.edu

PERFORMANCE OF SERIAL, MATCHED-FILTER PACKET ACQUISITION ALGORITHMS WITH A PREAMBLE-SEQUENCE ACCEPTANCE CRITERION

A Thesis
Presented to
the Graduate School of
Clemson University

In Partial Fulfillment
of the Requirements for the Degree
Master of Science
Electrical Engineering

---

by
Javier Schlömann
August 2007

Accepted by:
Dr. Daniel Noneaker, Committee Chair
Dr. Harlan Russell
Dr. Carl Baum

ABSTRACT

In this thesis, we consider two methods to improve the acquisition performance of a packet radio system that uses serial, matched-filter acquisition: an adaptive acquisition threshold, and an acceptance criterion for the system's preamble sequences. Each packet transmission includes a fixed-length acquisition preamble, and the preamble sequence used in packet transmissions is changed at predetermined times based on a sequence-generation algorithm. It is shown that acquisition performance depends largely on the sidelobe energy of the preamble sequence, the acquisition threshold, and the signal-to-noise ratio.

The first method uses a threshold-scaling technique to account for the variation in the signal-to-noise ratio. The second method employs preamble sequence-selection based on the preamble's sidelobe energy to reject sequences which are predicted to yield poor acquisition performance. The two techniques are considered individually and in combination, and the range of system parameters for which each is beneficial is investigated. Performance is examined for an additive white Gaussian noise channel and a signal-to-noise ratio that is unknown *a priori* at the receiver.

DEDICATION

This thesis is dedicated to my parents, Manfred and Lucia Schlömann, for their love and support.

ACKNOWLEDGMENTS

I would like to thank my advisor, Dr. Daniel Noneaker, for his insight and guidance during my research. His patience and dedication to my work have been vital throughout the preparation of this thesis. I would also like to thank Dr. Harlan Russell and Dr. Carl Baum for serving on my thesis committee.

In addition, I would like to thank all of the wonderful people who have been a part of my Clemson experience. I am extremely grateful to all of them for the memories which bring a smile to my face whenever I think of Clemson.

TABLE OF CONTENTS

Page

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1
INTRODUCTION

1.1 Packet Radio Systems

Packet radio systems are versatile and beneficial for use in both commercial and military applications. Random-access packet-based communications can provide advantages with respect to the efficiency of network resource utilization over circuit-switched communications for some cellular network services. Additionally, packet-based transmissions are the natural choice for robust communications in most circumstances in which access to a fixed communications infrastructure is unavailable, such as for a tactical radio network for the military or for disaster-relief operations.

In order to successfully receive a packet, the receiver must recognize the presence of an arriving packet, achieve a timing alignment with the incoming packet, and then proceed to despread the data-bearing portion of the packet. Despreading is achieved through the multiplication of the incoming packet with a locally generated replica of the transmission's spreading sequence synchronized with the spreading sequence of the incoming transmission. Timing alignment is usually realized through two stages: an initial coarse alignment is achieved in the *acquisition stage*, and a subsequent fine-tuning of the timing alignment is performed in the *tracking stage*.

The *a priori* uncertainty in the time of arrival of the next packet in a random-access communication system may be large, and consequently the performance of the acquisition stage of the receiver is often the limiting factor in the link-level performance of the radio network [1]. Thus, the design and implementation of the acquisition technique in a packet radio network is of particular interest when designing a system to meet certain acceptable performance criteria. This thesis addresses the design and performance of the acquisition stage.

1.2 Serial, Matched-Filter Packet Acquisition

Each packet transmission begins with a fixed-length preamble waveform defined by a preamble sequence with no superimposed data modulation, where the preamble sequence is known *a priori* at the receiver. For many systems, most notably for tactical radio systems deployed by the military, the preamble sequence used in packet transmissions must be changed frequently for security or other reasons. In this thesis, we consider a packet radio system in which the same preamble is used for all packet transmissions within a predetermined time interval known as a *time epoch*, but the preamble is changed at the start of each new epoch based on a sequence-generation algorithm used in common by the transmitting radio and the receiving radio.

The receiver we consider employs a serial, threshold-based acquisition algorithm using a filter matched to the preamble waveform. It is shown in

[2] that the intermediate-frequency (IF) filter and the subsequent automatic gain-control (AGC) system in a heterodyne receiver together have an effect on the performance of serial, matched-filter acquisition that is detrimental to the desired acquisition behavior in some surprising and significant respects. The effect of both components of the receiver are accounted for in the analysis in this thesis.

It is shown in [3] that a simple adaptive algorithm for setting the acquisition threshold can mitigate the phenomenon in a way that improves the performance of serial acquisition substantially in comparison with fixed-threshold acquisition for direct-sequence (DS) spread-spectrum communications. The windowing-based algorithm provides greater robustness with respect to variation in the strength of received packet transmissions as well as with respect to variation in the characteristics of the IF filter and the AGC system.

In [4], we consider a different approach to improving the performance of serial, matched-filter acquisition. In that approach, the preamble sequence for each epoch is determined by applying a sequence-acceptance criterion to the output of the pseudo-random preamble sequence generator. The investigation addresses the tradeoff between sequence selectivity and fixed-threshold acquisition performance on the one hand and the size of the pool of acceptable sequences and the search time to find an acceptable sequence on the other hand. It is shown that a properly designed selection criterion can be used to obtain a significant improvement in the average acquisition performance.

Moreover, this can be achieved while retaining a large pool of acceptable sequences and a reasonable average search time with simple sequence-generation algorithms.

In this thesis, we consider the performance of the sequence-selection technique of [4] and the adaptive-threshold acquisition algorithm of [3] both separately and when used together. The impact of each approach on acquisition performance is considered for short preamble sequences, such as might be employed in a narrowband communication system, for intermediate-length preamble sequences, and for long preamble sequences, such as might be employed in a DS spread-spectrum communication system.

1.3 Organization of Thesis

The thesis is organized as follows. In Chapter 2, we describe the basic communication system. Here we present the transmission format, the architecture of the receiver, and the algorithm used for packet acquisition. Chapter 3 describes the characterization of the test statistics as well as the evaluation of acquisition performance. The performance of the communication system using fixed-threshold acquisition is examined in Chapter 4. Optimization of the fixed threshold is described and an analysis of the performance with the min-max optimal fixed threshold is presented. Furthermore, a sequence-selection algorithm is introduced as one technique suitable for improving acquisition performance in some packet-radio networks and its performance is subsequently evaluated. Chapter 5 presents another technique to improve acquisition performance in some packet-based communication systems. Here we present the adaptive threshold algorithm and demonstrate how continuous adaptation of the acquisition threshold can be used to improve acquisition performance. Its use in conjunction with the sequence-selection criterion is addressed in Chapter 7 for different length preamble sequences. Conclusions are given in Chapter 8.

CHAPTER 2
SYSTEM DESCRIPTION

In this thesis, we consider a packet radio network in which time is divided into *epochs* with a duration much greater than the transmission time of a packet. The beginning and ending times and the identity of each epoch are known to all radios in the network. Each packet transmission in the network includes an acquisition preamble based on a sequence that depends on the identifier of the current epoch, perhaps along with characteristics of the transmission. Possible characteristics include the identity of the transmitting radio, the identity of the receiving radio, and the nature of the transmission such as the type of application traffic it contains or the function of the transmission within the radio-link protocol. Thus for a given transmitter-receiver pair and a given type of transmission, a fixed preamble sequence is used for all packet transmissions within a given epoch and a new preamble sequence is generated at the start of each epoch based on a pseudo-random algorithm.

The system definition is sufficiently general to encompass networks that use a common preamble sequence for all transmissions in the network within an epoch, networks that use fixed transmitter-directed preamble sequences [5] within an epoch, and networks that use fixed receiver-directed preamble sequences [5] within an epoch, among others. All radios in the network use the same sequence-generation algorithm(s) so that a consistent understanding of the preamble-sequence assignment(s) for the current time epoch exists at all radios. For simplicity, the description in the rest of the thesis is based on a network in which a common preamble is used by all transmissions in the network within a given epoch, though the analysis and conclusions are equally applicable in the general case. We focus on the acquisition performance of a single link.

2.1 Transmitted Signal and Received Signal

The transmitted signal for each packet consists of an acquisition preamble followed by the data-bearing signal of the packet. The baseband-equivalent, complex-valued preamble signal is

$$\tilde{s}(t) = \sqrt{P}\, c(t),$$

where P is the power in the transmitted signal during the preamble and $c(t)$ is the preamble waveform. The preamble waveform is given by

$$c(t) = \sum_{i=0}^{M-1} a_i \psi(t - iT_c),$$

where M is the preamble sequence length, $\psi(t)$ is the pulse waveform, and T_c is the pulse duration. The preamble sequence $\{a_i\}$ is a complex-valued quaternary sequence [6] taken from the values $\{e^{j\frac{\pi}{4}}, e^{j\frac{3\pi}{4}}, e^{j\frac{5\pi}{4}}, e^{j\frac{7\pi}{4}}\}$. Thus the radio-frequency (RF) transmission employs quadriphase-shift keyed modulation.

The preamble sequence for a given time epoch is randomly generated from among all length-M quaternary sequences that satisfy a selection criterion. (In some of the systems considered in this thesis, no selection criterion is enforced so that any length-M quaternary sequence may be generated.) This generalizes the *random sequence model* of [7] to account for the use of a sequence-selection criterion. A well-designed pseudo-random sequence generator using the selection criterion should have statistical properties similar to that of the random sequence model with the same selection criterion.

The pulse waveform is time-limited to $[0, T_c)$ and has unit average power. If narrowband modulation is used for the packet transmission, the pulse duration corresponds to the duration of a channel-symbol in the data-bearing signal. If instead DS spread-spectrum modulation is used, the pulse duration corresponds to the duration of a chip in the spreading waveform of the data-bearing signal.

The transmitted signal arrives at the receiver over an additive white Gaussian noise (AWGN) channel. Without loss of generality, we assume that the signal is neither attenuated nor delayed by the channel. Thus, the baseband-equivalent received signal is given by

$$\tilde{r}(t) = \tilde{s}(t) + \tilde{n}(t),$$

where the complex-valued noise process, $\tilde{n}(t)$, has two-sided power spectral density N_0.

2.2 Receiver Architecture

2.2.1 Heterodyne Receiver

The communication system considered in this thesis includes a heterodyne receiver [8], though the model and analysis used in the thesis are readily adapted to a system using a direct-conversion (i.e., zero-IF) receiver [8]. A heterodyne receiver consists of an RF-IF front end, one or more IF down-conversion stages, and baseband signal processing. The RF-IF front end consists of an antenna, a low-noise amplifier, an image-rejection filter, and an RF-to-IF down-conversion mixer. The most common heterodyne receiver architecture employs a single IF down-conversion stage which typically consists of an IF noise-limiting filter (henceforth referred to as the *IF filter*), an AGC system, and a down-converter from the IF to baseband. The receiver architecture considered in this thesis employs only a single IF stage, though the model and analysis used in the thesis can be adapted to a system with multiple IF stages. The structure of a typical single-IF heterodyne receiver is illustrated in Fig. 2.1.

The choice of the intermediate frequency in the single-IF heterodyne receiver determines a tradeoff between the stringency of requirements imposed on the image-rejection filter, the RF-to-IF mixer, and the IF filter. In particular, a low-IF receiver design [8] allows for a cost-effective IF filter with a bandwidth not much greater than the bandwidth of the desired signal in

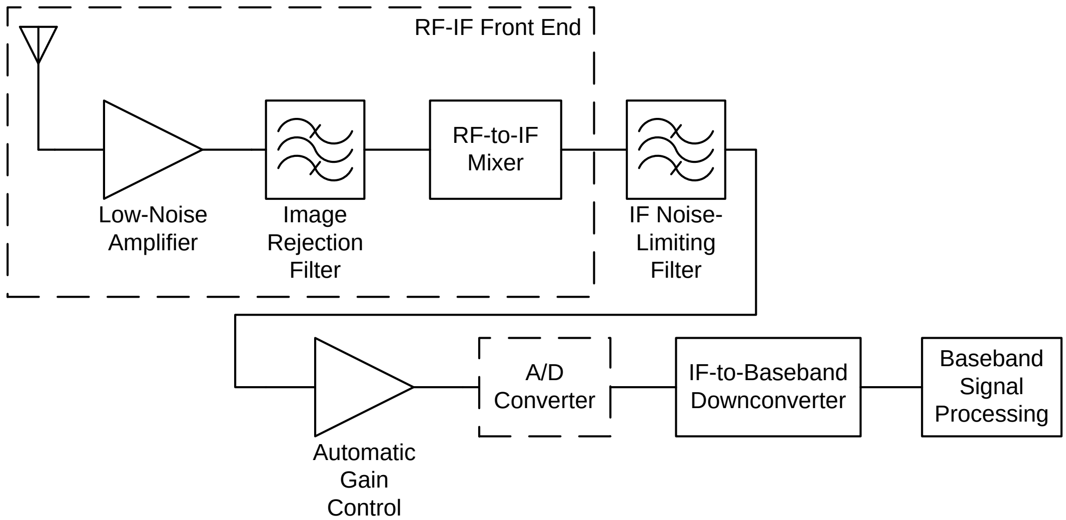

Figure 2.1: Typical structure of a single-IF heterodyne receiver.

many instances [9], whereas a high-IF receiver design [8] typically requires an IF filter with a bandwidth that is a larger multiple of the signal bandwidth [9].

Modern communication receivers perform baseband signal processing digitally, necessitating the use of an analog-to-digital (A/D) converter. The location of the AGC system and the A/D converter in the sequence of receiver stages differs for different receiver designs, though their location as shown in Fig. 2.1 is used in most single-IF heterodyne receivers. The model of the receiver used in this thesis assumes that the AGC system is located as shown in the figure.

2.2.2 Baseband-Equivalent Receiver

A well-designed RF-IF front end results in negligible distortion of the desired signal beyond the additive white Gaussian noise introduced in the antenna and low-noise amplifier. Consequently, its effect on the baseband-equivalent received signal can be represented by an AWGN source. Moreover, quantization error due to the A/D conversion has only a secondary effect on the performance of many digital communication systems. In such instances, the presence of the A/D converter need not be reflected in a model of the receiver. We employ both of these approximations in this thesis, so that the receiver considered in the thesis is shown in Fig. 2.1 with the subsystems surrounded by dashed lines omitted.

The model of the heterodyne receiver we consider thus includes an AWGN source, a bandpass IF filter, and an AGC system followed by an acquisition stage as shown in Fig. 2.2. (The model can be adapted for a direct-downconversion receiver, which is also known as a zero-IF receiver [8].) The baseband equivalent of the IF filter's impulse response is denoted by $\tilde{h}(t)$, and the baseband-equivalent received signal $\tilde{r}(t)$ is approximated as undistorted prior to IF filtering.

A well-designed AGC system varies its average output power minimally over a wide range for the input power. This behavior is approximated by an AGC system that responds instantly to a step change in the average steady-state input power and maintains a constant average steady-state output power.

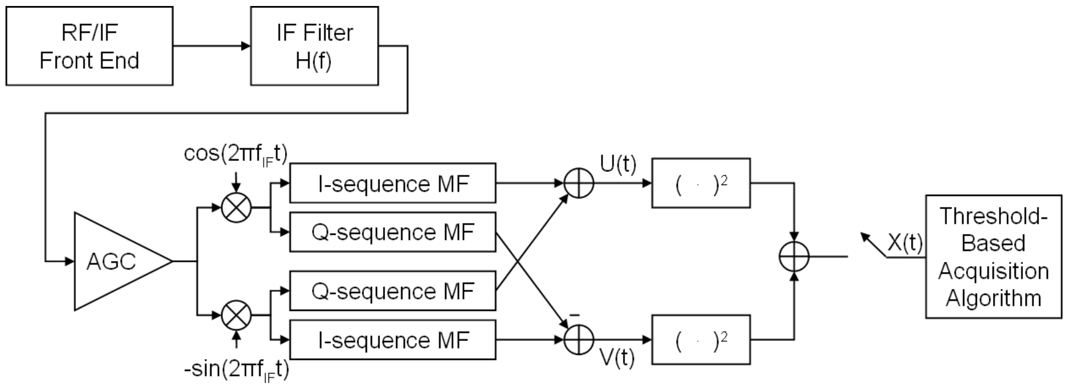

Figure 2.2: Block diagram of the receiver architecture.

The baseband equivalent of the AGC output signal is thus given by

$$\hat{r}(t) = \sqrt{\alpha(t)}(\tilde{r} * \tilde{h})(t),$$

where $1/\alpha(t)$ is the average steady-state input power to the AGC system at time t. The quantity $\alpha(t)$ depends on the input signal power, the noise power spectral density, the pulse waveform, and the IF filter's impulse response. It is shown in [2] that it can be expressed as

$$\alpha(t) = \begin{cases} \alpha_0 = \left(\gamma_s \beta \frac{N_0}{T_c}\right)^{-1}, & t < 0 \\ \alpha_1 = \gamma_s^{-1}\left(P + \beta \frac{N_0}{T_c}\right)^{-1}, & 0 \leq t \leq MT_c \end{cases} \tag{2.1}$$

where

$$\begin{aligned} \gamma_s &= \int_{-\infty}^{+\infty} \frac{|\Psi(f)|^2}{T_c} |\tilde{H}(f)|^2 \mathrm{d}f, \\ \gamma_n &= T_c \int_{-\infty}^{+\infty} |\tilde{H}(f)|^2 \mathrm{d}f, \end{aligned}$$

and

$$\beta = \frac{\gamma_n}{\gamma_s}.$$

The function $\Psi(f)$ is the Fourier transform of the pulse waveform $\psi(t)$, $\tilde{H}(f)$ is the baseband-equivalent of the IF filter's frequency response, and the parameter β is a measure of the signal-normalized bandwidth of the IF filter. The parameter $\beta \geq 1$ in all instances [2]. Ideal Nyquist signaling with a matched IF filter would result in a $\beta = 1$. Typical IF filters of low-IF heterodyne receivers result in values of β in the range of two to three, whereas those of high-IF heterodyne receivers result in larger values of β. For each system considered in the examples in this thesis, the chip waveform and IF filter result in a value of $\beta = 3.0$.

2.3 Acquisition Algorithm

A receiver awaiting the arrival of a packet is in *acquisition mode*. Acquisition of an arriving packet is achieved by serial application of a threshold-based test to non-coherent square-law statistics. In-phase and quadrature branches of the receiver each contain two filters, denoted by $g_1(t)$ and $g_2(t)$, where $g(t) = g_1(t) + jg_2(t)$ is matched to the baseband-equivalent (complex) preamble waveform. The outputs of the filters are sampled periodically, and the two samples at time t are used to form a test statistic

$$X(t) = [U(t)]^2 + [V(t)]^2 \,,$$

where

$$U(t) = \mathcal{R}e\left\{(\hat{r} * g)(t)\right\}$$

and

$$V(t) = \mathcal{I}m\left\{(\hat{r} * g)(t)\right\}.$$

Each test statistic generated in acquisition mode is compared with an *acquisition threshold* η. If $X(t) > \eta$, a *hit* is declared, which provides a tentative indication that the end of a packet's preamble has arrived at time t.

Suppose the receiver enters acquisition mode and a packet arrives at some subsequent time. If no hit occurs in the interval prior to the packet's arrival, the interval during the arrival of the packet's preamble, or within the interval afterward corresponding to the pull-in range of the receiver's pulse-delay tracking algorithm [10], a *miss* is said to occur and the receiver fails to acquire the packet. If instead a hit occurs at a time that precedes the arrival of the end of the preamble by more than the pull-in range, a *false alarm* occurs. If a hit occurs at a time that differs from the arrival of the end of the preamble by an amount less than the pull-in range, however, successful acquisition occurs.

When a hit occurs, the receiver switches from acquisition mode to *verification mode*. During verification mode, the receiver employs a *verification test* which is designed to determine whether a hit corresponds to successful acquisition or a false alarm. The duration of the verification test is the *verification interval*. If successful acquisition is verified for a hit at time t, the timing information is used to synchronize a locally generated reference signal with the received signal for subsequent demodulation of the data symbols in the packet. The receiver then switches from verification mode to *data-detection mode*. If

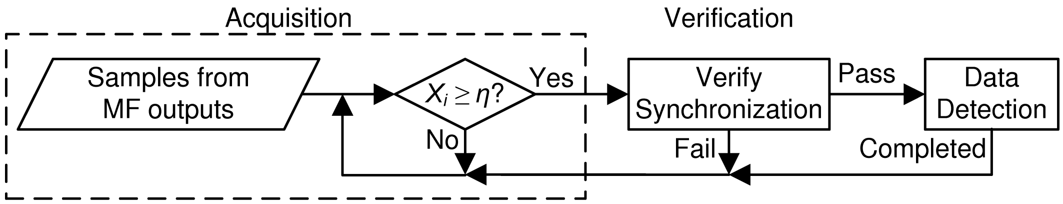

Figure 2.3: Modes of the receiver.

instead a false alarm is identified during verification mode, the receiver returns to the acquisition mode and once again awaits the arrival of a packet. (In some receiver designs, tentative data detection may begin concurrently with verification and terminate if verification fails.) The relationship among the three modes of the receiver is shown in Fig. 2.3.

The receiver design considered in this thesis does not utilize subsequent outputs of the preamble-matched filter while it is in verification mode. Thus if the occurrence of a false alarm is such that the receiver is in verification mode when the end of a packet's preamble arrives, the false alarm results in a failure to acquire the packet. In all that follows, we thus use the term "false alarm" to mean a false alarm that results in a failed packet acquisition. Examples of a miss, a false alarm, and a successful packet acquisition are shown in Figs. 2.4-2.6 for a receiver using a fixed acquisition threshold. Each of the three illustrates a sequence of periodically generated samples used as test statistics, the time of arrival of the end of a packet's preamble (i.e., the time at which the packet may be successfully acquired), and the receiver mode that is active at each moment.

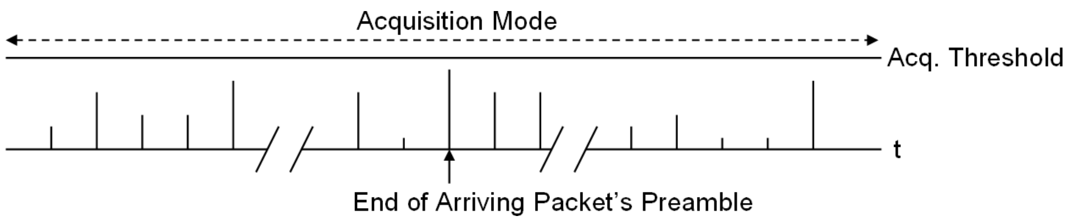

Figure 2.4: Depiction of a miss.

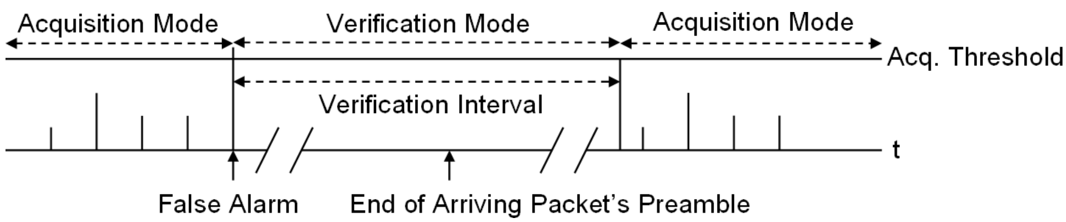

Figure 2.5: Depiction of a false alarm.

Depending on the design of the verification test, the verification interval can be either constant or variable. Since the design of the verification test is not a focus of this thesis, we consider only a constant verification interval of Q times the pulse duration. When considering a packet that arrives beginning at time $t = 0$, we assume that the receiver is in the acquisition mode at time $t = (M - Q)T_c$ and that no other packet transmissions are present during the interval $[(M-Q)T_c, MT_c]$. (We extend the time interval to which the latter assumption applies when considering a receiver using the adaptive-threshold algorithm, as discussed in Section 5.2.) A well-designed verification test provides a correct result with a very high probability, and we employ the approximation that the verification test always gives a correct result.

In practice, a sampling period of $T_c/2$ or less is used to compensate for the

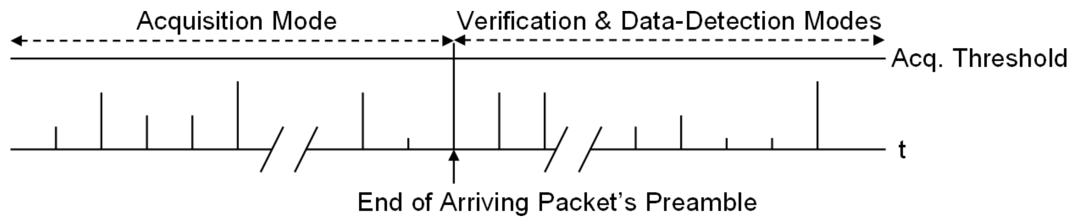

Figure 2.6: Depiction of successful acquisition.

unknown *a priori* sample-timing error relative to the received pulse waveform. The assumption of pulse-rate sampling and pulse-level synchronism at the receiver results in an accurate approximation to the acquisition performance with over-sampling acquisition and a tracking delay-locked loop for an arbitrary sample-timing error [10], however, and thus we employ that assumption in this thesis. Consequently, we assume the filter outputs are sampled by the receiver at times $t = iT_c$, where i is an integer, with the filter outputs denoted U_i and V_i and the corresponding test statistic is given by

$$X_i = U_i^2 + V_i^2. \tag{2.2}$$

We also assume that the pull-in range of the receiver's pulse-delay tracking loop is less than T_c, so that successful acquisition occurs only if test statistic X_M results in a hit.

CHAPTER 3

EVALUATION OF ACQUISITION PERFORMANCE

The performance of each acquisition algorithm considered in this thesis is characterized in terms of the average acquisition performance over all time epochs in the system. Thus it is the average of the acquisition performance over those quaternary preamble sequences of a given length M that satisfy the sequence-acceptance criterion employed by the system. Equivalently, it is the expected acquisition performance with a random, uniformly distributed quaternary preamble sequence subjected to the sequence-acceptance criterion. For those systems in which the preamble sequence is not constrained by an acceptance criterion, it is simply the expected performance with a random preamble sequence.

Any reference to a measure of acquisition performance in this thesis implies the average or expected performance, except where otherwise explicitly noted. All numerical examples concern the average performance. In a few places in this chapter and the next one, there is analysis and discussion of the acquisition performance for a given preamble sequence. Such instances are identified explicitly.

The performance measure of principal interest is the probability of not acquiring,

$$\mathrm{P}_{\mathrm{nacq}} = 1 - \mathrm{P}(X_M \geq \eta, X_{M-1} < \eta, \ldots, X_{M-Q} < \eta).$$

Two ancillary measures of performance are also useful in providing insight into the behavior of an acquisition algorithm. They are the probability of miss,

$$\mathrm{P}_{\mathrm{miss}} = \mathrm{P}(X_M < \eta, \cdots, X_{M-Q} < \eta),$$

and the probability of false alarm,

$$\mathrm{P}_{\mathrm{fa}} = 1 - \mathrm{P}(X_{M-1} < \eta, \cdots, X_{M-Q} < \eta),$$

which are related to the probability of not acquiring by

$$\mathrm{P}_{\mathrm{nacq}} = \mathrm{P}_{\mathrm{miss}} + \mathrm{P}_{\mathrm{fa}}. \tag{3.1}$$

(Note that a different definition of $\mathrm{P}_{\mathrm{miss}}$ is used in [2], and (3.1) does not hold for that definition.) Performance is evaluated using the approximation that the baseband-equivalent IF filter introduces no distortion in the pulse waveform (that is, $\psi(t) = (\psi * \tilde{h})(t)$), which is an accurate approximation for most receiver designs.

Conditioned on the preamble sequence $\{a_0, \ldots, a_{(M-1)}\}$, the performance of a given acquisition algorithm depends on the *preamble signal-to-noise ratio* at the receiver and the *aperiodic autocorrelation function* [6] of the preamble sequence. The preamble signal-to-noise ratio is given by

$$\mathrm{SNR} = \frac{MPT_c}{N_0},$$

and the aperiodic autocorrelation function is the conjugate-symmetric discretetime function

$$C_a(k) = \begin{cases} \sum_{j=0}^{M-1-k} a_j a_{j+k}^*, & 0 \le k \le M-1 \\ 0, & |k| \ge M \end{cases} . \tag{3.2}$$

Conditioned on the preamble sequence, the test statistics $\{X_i\}$ in (2.2) are quadratic forms in the jointly Gaussian random variables $\{(U_i, V_i)\}$ [11] and are thus characterized by the first and second moments of $\{(U_i, V_i)\}$ [3]. It is shown in [3] that conditioned on the preamble sequence,

$$|E[U_i + jV_i]| = \begin{cases} 0, & i \le 0 \\ \frac{T_c|C_a(M-i)|}{\sqrt{\gamma_s\left(1+\frac{M\beta}{\text{SNR}}\right)}}, & 1 \le i < M \\ \frac{T_c M}{\sqrt{\gamma_s\left(1+\frac{M\beta}{\text{SNR}}\right)}}, & i = M \end{cases} , \tag{3.3}$$

$$\text{Var}(U_i) = \text{Var}(V_i) = \begin{cases} \frac{MT_c^2}{2\gamma_s\beta}, & i \le 0 \\ \frac{(M-i)T_c^2}{2\gamma_s\beta} + \frac{MiT_c^2}{2\gamma_s(\text{SNR}+M\beta)}, & 1 \le i \le M \end{cases} , \tag{3.4}$$

and

$$\text{Cov}(U_i, V_i) = 0$$

for all i. The two statistics within each of the pairs (U_i, U_j), (V_i, V_j), and (U_i, V_j) are dependent in general for $i \neq j$, however, due to the correlation of their noise components. Thus the test statistics are also dependent in general, and neither the probability of false alarm nor the probability of not acquiring can be expressed exactly in a closed form.

The performance of a given acquisition algorithm can vary significantly among preamble sequences, depending on the aperiodic autocorrelation function of the preamble sequence [3]. For most preamble sequences, the off-peak magnitudes of the aperiodic autocorrelation function (i.e., $|C_a(k)|$ for $k \neq 0$) are sufficiently small that the approximation $\text{Cov}(U_i, U_j) \approx 0$ and $\text{Cov}(V_i, V_j) \approx 0$ for all $i \neq j$ is accurate. Furthermore, $\text{Cov}(U_i, V_j) \approx 0$ for all i and j regardless of $C_a(k)$. Thus the test statistics $\{X_i\}$ can be accurately approximated as independent non-central chi-square random variables. If a fixed threshold is used in the acquisition algorithm, the probability of not acquiring that results with the specified preamble sequence is given by a closed-form expression that involves the Marcum Q-function [12]. (This result is presented in [3], but the conditions of the approximation are not stated clearly there.) Even under the approximation, the acquisition performance depends heavily on the aperiodic autocorrelation function of the preamble sequence.

The performance of the acquisition algorithm averaged over all preamble sequences generated by the system depends on the sequence-generation algorithm that the system employs. If the sequence-generation algorithm is modeled as a random generator of quaternary preamble sequences with no constraints on the generated sequences, a simplifying approximation of the test statistics can be exploited as shown in [2]. For a random preamble sequence, the pairs of random variables $\{(U_i, V_i)\}$ are mutually uncorrelated with

$$|E\left[U_i + jV_i\right]| = \begin{cases} 0, & i \leq M-1 \\ \frac{T_c M}{\sqrt{\gamma_s \left(1+\frac{M\beta}{\text{SNR}}\right)}}, & i = M \end{cases} \tag{3.5}$$

and

$$\text{Var}(U_i) = \text{Var}(V_i) = \begin{cases} \frac{MT_c^2}{2\gamma_s\beta}, & i \leq 0 \\ \frac{MT_c^2}{2\gamma_s\beta} + \frac{iT_c^2(\beta-1)}{2\gamma_s\beta} \times \left(\frac{1}{1+\frac{\beta M}{\text{SNR}}}\right), & 1 \leq i < M \cdot \\ \frac{MT_c^2}{2\gamma_s\beta}\left[1 - \frac{1}{1+\frac{\beta M}{\text{SNR}}}\right], & i = M. \end{cases} \tag{3.6}$$

The joint distribution function of the collection of uncorrelated random variables $\{(U_i, V_i)\}$ can be approximated as jointly Gaussian. Under the approximation, the random variables are thus independent Gaussian random variables with first and second moments given by (3.5) and (3.6), respectively, and the test statistics are independent chi-square random variables. Furthermore, all the test statistics except X_M are *central* chi-square random variables. If a fixed threshold is used in the acquisition algorithm, the probability of not acquiring a random preamble sequence can be expressed in a simple closed form under the Gaussian approximation [2]. It is shown in [3] that the Gaussian approximation leads to extremely accurate results for the probability of not acquiring with a fixed threshold for preamble sequence lengths and verification intervals of practical interest.

While the closed-form expressions described above provide accurate approximations to the exact probability of not acquiring for the systems to which they are applicable, their applicability is not sufficiently broad to encompass most of the systems considered in this thesis. The sequence-selection algorithms considered in this thesis result in the acceptance of only a subset of all length-M quaternary sequences as valid preamble sequences, and a dynamic acquisition threshold is used in some of the systems considered. No accurate closed-form approximation to the probability of not acquiring is available that accounts for either of these factors. Consequently, all performance evaluation in this thesis employs Monte Carlo simulation of each system based on the jointly-Gaussian approximation of the joint distribution function of the test statistics.

CHAPTER 4

ACQUISITION USING A FIXED THRESHOLD

The simplest serial, matched-filter acquisition algorithm uses a fixed acquisition threshold, and it provides the most straightforward insights into the behavior of this class of acquisition algorithms. In this chapter, we consider the acquisition performance with a fixed acquisition threshold. We also discuss the underlying cause of a counter-intuitive characteristic of the system and illustrate its detrimental effect on the performance. This motivates a performance criterion that is particularly well suited to the behavior of the system, and we describe a previously developed method for determining the optimal acquisition threshold with respect to the performance criterion.

4.1 Performance of the Fixed-Threshold Acquisition Algorithm

Consider the packet radio system described in Chapter 2. The use of an IF filter with a bandwidth greater than the bandwidth of the desired signal results in an AGC system which responds in part to the noise power that falls within the passband of the IF filter but outside the passband of the IF-equivalent desired signal. Thus as the strength of the desired signal increases (and the signal-to-ratio increases proportionally), the decrease in the gain of the AGC system does not fully offset the increase in the input power within the signal's passband. Consequently, the total power at the output of the AGC system within the signal's passband increases as the signal-to-noise ratio increases. As seen from (2.1), this mismatch between the AGC system's behavior and its design objective of a constant output power becomes more pronounced as the relative bandwidth of the IF filter is increased (i.e, as β is increased).

This undesirable phenomenon is reflected in the distribution function of the test statistics of the acquisition algorithm as well. It is apparent that the parameters of the test statistic X_i which are given by (3.3) and (3.4) increase in value with an increasing signal-to-noise ratio for $1 \leq i \leq M$. Thus while the distribution function of the chi-square random variable X_i does not depend on the signal-to-noise ratio for $i < 0$, it increases stochastically with an increasing signal-to-noise ratio for $1 \leq i \leq M$. Moreover, the aperiodic autocorrelation function of each preamble sequence tends to decrease in magnitude with an increase in its argument so that the parameter given in (3.3) tends to increase in value with increasing i for $1 \leq i \leq M$. Consequently, the presence of non-zero side lobes in the aperiodic autocorrelation function amplifies the stochastic increase in the distribution function of each X_i with increasing signal-to-noise ratio so that it is more pronounced for a typical test statistic generated closer to the end of the received preamble than for one generated earlier.

Consider serial, matched-filter acquisition using a fixed acquisition threshold. Within the approximation that the test statistics $\{X_i\}$ are independent for a given preamble sequence, it follows that for any given preamble sequence, the probability of miss is a decreasing function of the signal-to-noise ratio whereas

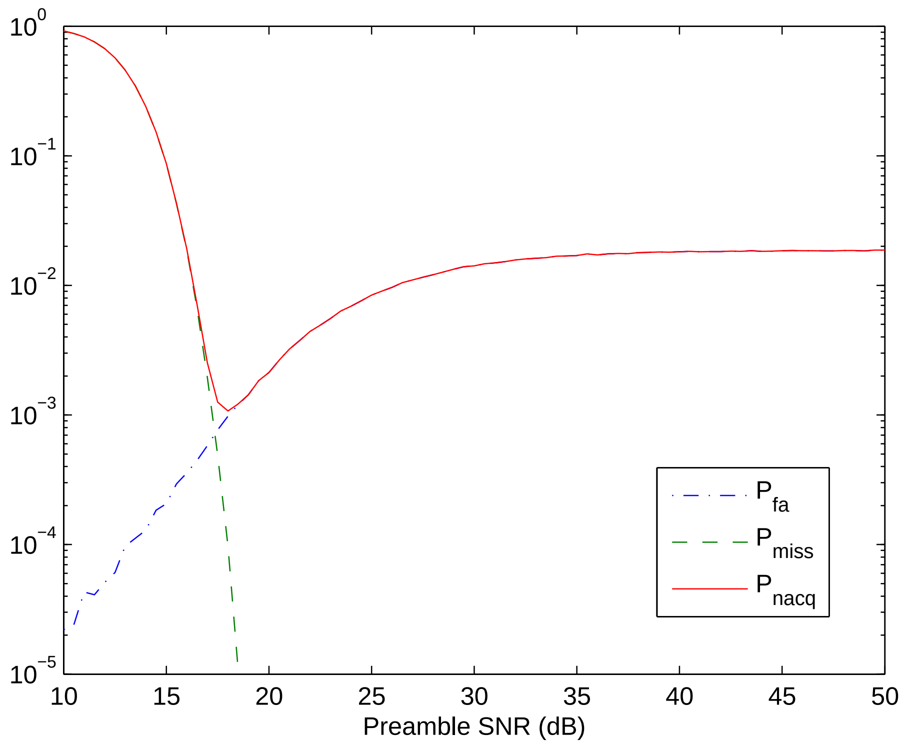

Figure 4.1: Performance for preamble sequences of length 26.

the probability of false alarm is an *increasing* function of the signal-to-noise ratio. The probability of not acquiring with that sequence is thus a *non-monotonic* function of the signal-to-noise ratio, as is the average probability of not acquiring over all preamble sequences [2]. This is illustrated in Fig. 4.1, which reproduces a result from [3]. The figure shows the probabilities of miss, false alarm, and not acquiring averaged over randomly generated preamble sequences of length 26 using a particular choice of the acquisition threshold. The verification interval is 65 times the pulse duration.

4.2 Optimization of the Threshold

It is desirable to select the acquisition threshold for a particular system so that it yields the best acquisition performance over the full range of operating conditions of interest for that system. For a system communicating over an AWGN channel, the operating conditions of interest are commonly specified in terms of a minimum signal-to-noise ratio above which acceptable performance must be achieved. This design objective is complicated by the fact that the probability of not acquiring does not decrease monotonically with an increasing signal-to-noise ratio, however.

Consider a system using a fixed acquisition threshold that is initially set to a certain value. If that threshold is decreased, the probability of miss will decrease. The probability of not acquiring with any given preamble sequence will thus decrease if the preamble signal-to-noise ratio is low enough that misses predominate over false alarms. Yet the probability of false alarm will increase and the probability of not acquiring will thus increase, if the preamble signal-to-noise ratio is high enough that false alarms predominate over misses. Conversely, an increase in the threshold will result in an increase in the probability of not acquiring for a low signal-to-noise ratio but a decrease in the probability of not acquiring for a high signal-to-noise ratio. Consequently, the worst-case performance over a range of values of the preamble signal-to-noise ratio occurs at a signal-to-noise ratio that depends on the threshold [2].

The effect of the threshold on the average acquisition performance over randomly generated sequences is illustrated in Fig. 4.2 for the same system considered in Fig. 4.1. Consider the design objective of achieving the smallest possible worst-case probability of not acquiring over all channels with a preamble signal-to-noise ratio of at least 16 dB. If the acquisition threshold is low, the probability of false alarm is high and the worst-case probability of not acquiring occurs at a high preamble signal-to-noise ratio. If the acquisition threshold is high, the probability of miss is high and the worst-case probability of not acquiring occurs at a low preamble signal-to-noise ratio. A properly chosen intermediate value of the threshold balances these two extremes and results in better worst-case performance over the range of channels of interest.

The result of Fig. 4.2 motivates an approach to optimization of the acquisition threshold using the threshold-selection criterion presented in [2]. The performance measure used in the optimization is the average probability of not acquiring over all randomly generated preamble sequences of a given length (in some instances constrained to a sequence-acceptance criterion). Suppose that $\mathrm{SNR}_{\mathrm{min}}$ denotes the smallest preamble signal-to-noise ratio of interest. The threshold that minimizes the worst-case probability of not acquiring over all $\mathrm{SNR} \geq \mathrm{SNR}_{\mathrm{min}}$ is referred to as the *min-max optimal threshold*, and it is said to satisfy the *min-max criterion* [3]. The result of Fig. 4.2 also provides an illustration of threshold optimization using the min-max criterion. The intermediate threshold value that is considered is in fact the min-max optimal threshold for SNR_{min} =16 dB, and it results in a worst-case probability of not acquiring that is lower than that achieved with either of the other two thresholds considered.

It is seen in Fig. 4.2 that the min-max optimal threshold results in a probability of not acquiring for $\text{SNR} = \text{SNR}_{\text{min}}$ that is identical to the limiting probability of not acquiring as the signal-to-noise ratio approaches infinity. Indeed, it is shown in [2] that this is always true for the min-max optimal acquisition performance using a fixed threshold with a random preamble sequence. It is also shown in [2] that for a random preamble sequence, the probability of miss is a strictly decreasing function of the signal-to-noise ratio and the probability of false alarm is a strictly increasing function of the signal-to-noise ratio. Consequently, the min-max optimal threshold is unique for a given system and value of SNR_{min}. We exploit these facts in an algorithm for finding the min-max optimal fixed threshold. The algorithm is illustrated in Fig. 4.3.

In each example in this thesis, we consider the acquisition performance when the acquisition threshold is chosen according to the min-max criterion with $\text{SNR}_{\text{min}} = 16$ dB. Note that the channel-symbol signal-to-noise ratio in the data portion of the packet is typically much lower than the preamble signal-to-noise ratio. For example, for narrowband modulation and $M = 26$, a preamble signal-to-noise ratio of 16 dB corresponds to a channel-symbol signal-to-noise ratio of about 2 dB. The choice of SNR_{min} for the examples is

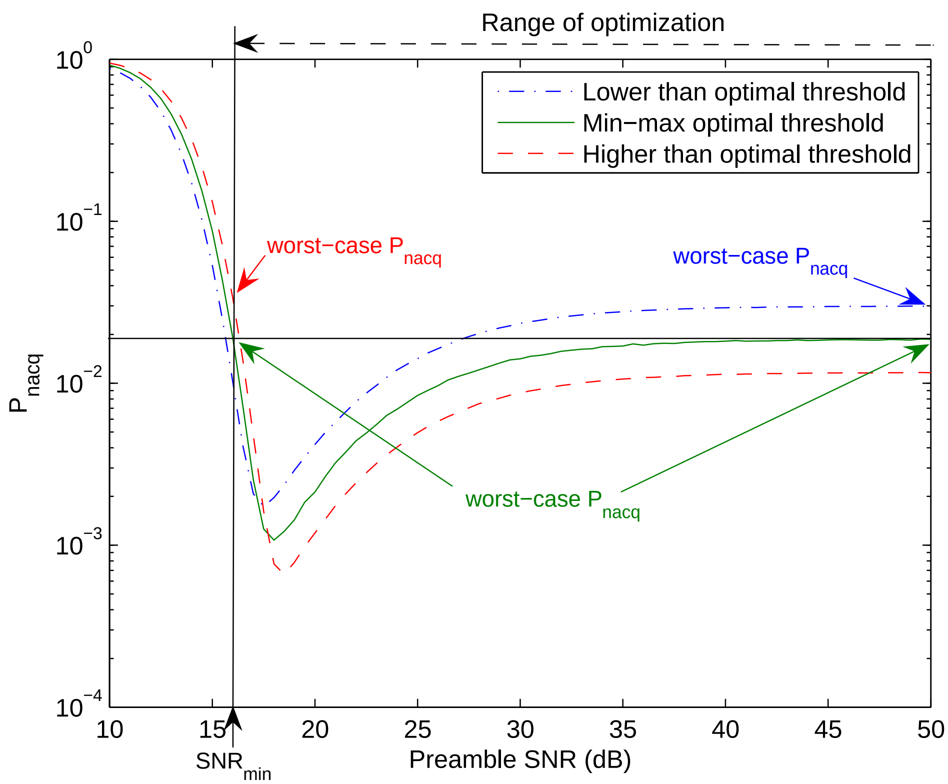

Figure 4.2: Performance of fixed-threshold acquisition for three values of the threshold and preamble sequences of length 26.

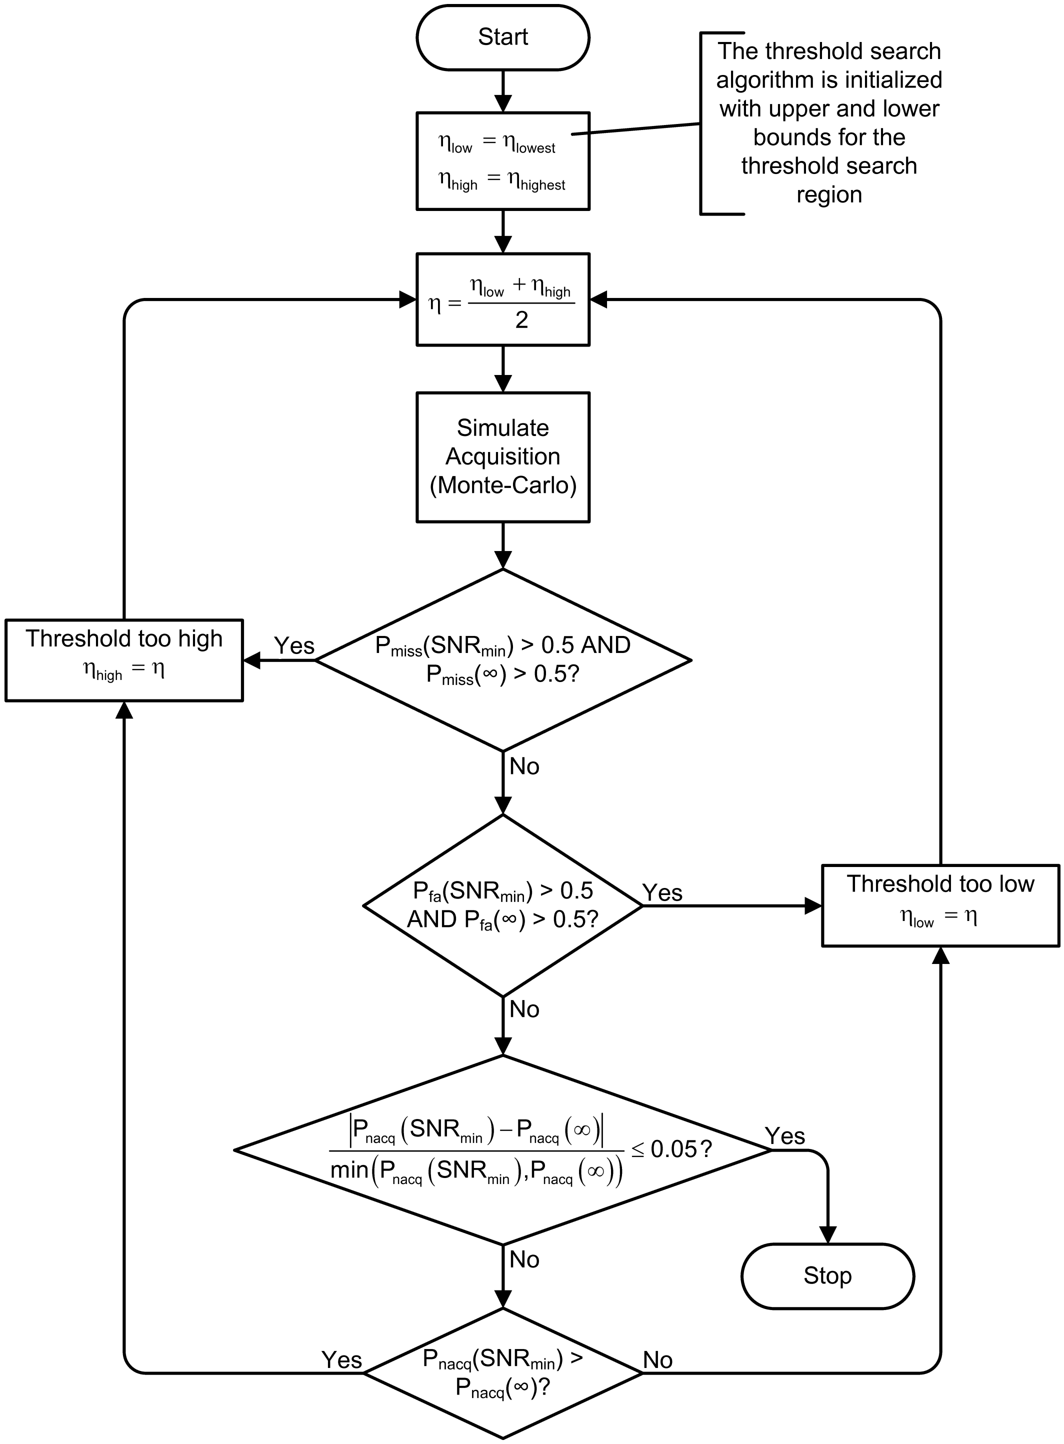

Figure 4.3: Algorithm for finding the min-max optimal fixed threshold.

appropriate for many practical systems.

4.3 Performance with the Min-Max Optimal Threshold

In this section, we illustrate the acquisition performance that results using the min-max optimal fixed threshold by considering systems using preamble sequences of three different lengths: $M = 26$, $M = 100$, and $M = 400$. The corresponding systems are said to employ *short preamble sequences*, *intermediate-length preamble sequences*, and *long preamble sequences*, respectively. A verification interval of $Q = 65$ times the pulse duration is considered with the short preamble sequences, whereas verification intervals of 250 times the pulse duration and 1,000 times the pulse duration are considered with the intermediate-length preamble sequences and the long preamble sequences, respectively. The same three combinations of the sequence length and the verification interval are used in all examples in subsequent chapters, except in Section 7.4.

The verification interval QT_c is 2.5 times the preamble duration MT_c for all three systems. Thus the systems can be viewed as all using the same preamble duration and having the same verification interval but employing three different pulse durations. This in turn can be viewed as corresponding to three different choices of DS spread-spectrum processing gain for the data symbols, ranging from narrowband communications or a small processing gain if $M = 26$ to an approximately sixteen-fold greater processing gain if $M = 400$.

The probability of not acquiring is shown in Fig. 4.1, Fig. 4.4, and Fig. 4.5

for the systems using short preamble sequences, intermediate-length preamble sequences, and long preamble sequences, respectively. The min-max probability of not acquiring is approximately 2×10^{-2} for each system, and it occurs for both a preamble signal-to-noise ratio of 16 dB and for a very large signal-to-noise ratio. The performance of the three systems differs considerably if the signal-to-noise ratio falls between these two extremes, however. If the preamble signal-to-noise ratio is 18 dB, for example, the probability of not acquiring is approximately 10^{-3} if $M = 26$ but less than 10^{-4} if $M = 100$ or $M = 400$. The non-monotonicity of performance as a function of the signal-to-noise ratio becomes more pronounced as the preamble sequence length is increased. The min-max optimal fixed threshold also differs significantly for the different system. The optimal threshold is $\eta = 0.2021 \left(\frac{M^2 T_c^2}{\gamma_s}\right)$ for the system with short preamble sequences, it is $\eta = 0.06812 \left(\frac{M^2 T_c^2}{\gamma_s}\right)$ for the system with intermediate-length preamble sequences, and it is $\eta = 0.01898 \left(\frac{M^2 T_c^2}{\gamma_s}\right)$ for the system with long preamble sequences.

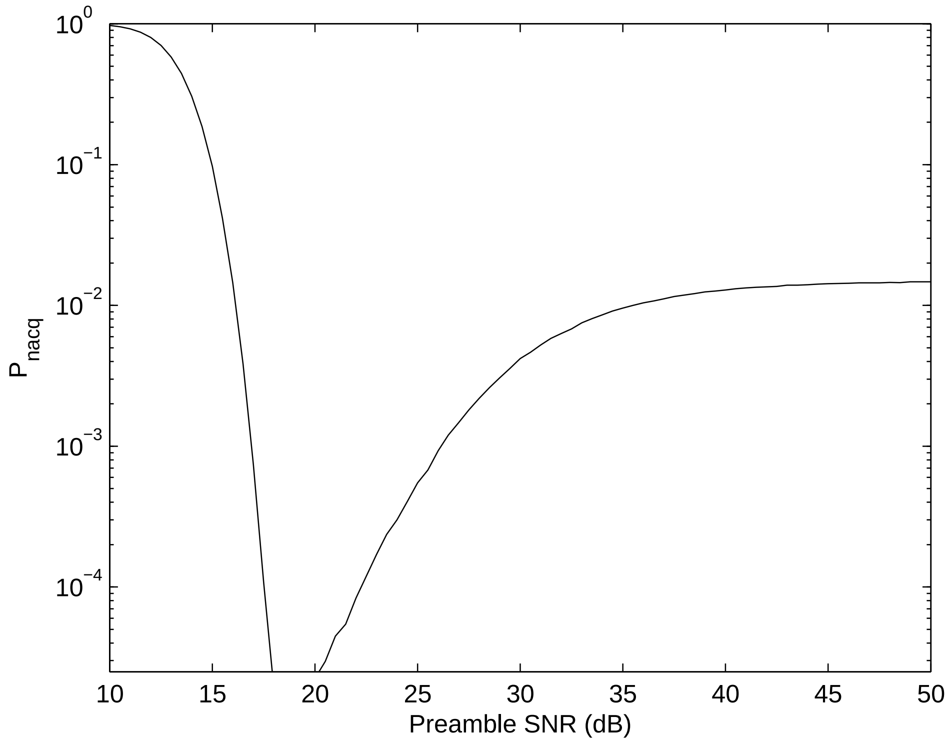

Figure 4.4: Min-max-optimal performance for preamble sequences of length 100.

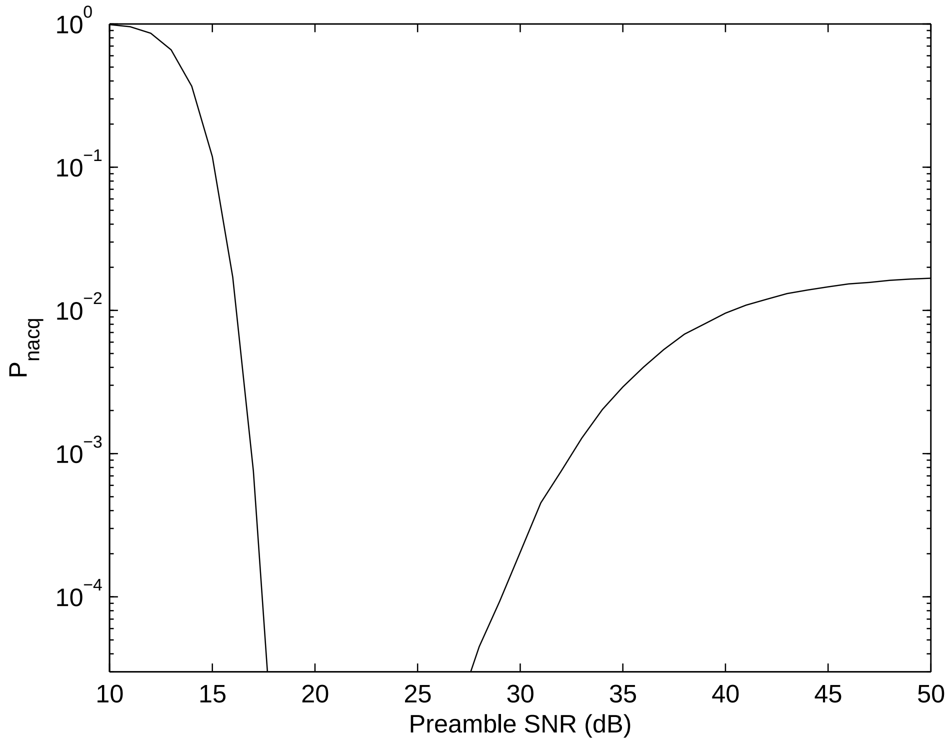

Figure 4.5: Min-max-optimal performance for preamble sequences of length 400.

CHAPTER 5

ACQUISITION USING AN ADAPTIVE THRESHOLD

It is shown in the previous chapter that the average probability of not acquiring over a randomly selected preamble sequence is a non-monotonic function of the preamble signal-to-noise ratio if serial, matched-filter acquisition is used with a fixed acquisition threshold. A simple modification to the acquisition algorithm is introduced in [3] which mitigates this undesirable behavior and results in improved acquisition performance. The modified algorithm employs a threshold that is adapted for each test statistic. In this chapter we describe the adaptive-threshold technique of [3], and we investigate the performance of the acquisition algorithm with preamble sequences of different lengths.

5.1 Adaptive-Threshold Algorithm

The disappointing performance obtained with a fixed acquisition threshold is a consequence of the stochastic increase with increasing signal-to-noise ratio for those test statistics which result in a false alarm after the arrival of the beginning of the preamble. Thus the increase in false alarms with increasing signal-to-noise ratio can be mitigated if the increase in the statistics can be estimated and some compensation can be employed. The adaptive-threshold algorithm uses this approach based on an estimate of the expected value of each test statistic.

The estimate formed for the test statistic X_i is the un-weighted average of the W most recently generated test statistics, where W is the *window size*. The estimate is given by

$$S_i = \frac{\sum_{k=i-W}^{i-1} X_k}{W}$$

where S_i is referred to as the *scaling factor*. The scaling factor is used to determine the acquisition threshold η_i against which the test statistic X_i is compared. The ith threshold is given by

$$\eta_i = \eta \frac{S_i}{\left(\frac{MT_c^2}{\gamma_s \beta}\right)}$$

where the constant η is the *nominal threshold*.

5.2 Performance of the Adaptive-Threshold Acquisition Algorithm

The effectiveness of the scaling factor S_i as an estimator of $E[X_i]$ differs for different test statistics. For test statistics generated prior to the arrival of the preamble, the scaling factor is an unbiased estimator of of $E[X_i]$. (That is, $E[S_i] = E[X_i]$ for $i \leq 0$.) For test statistics generated between the arrival of the beginning of the preamble and the arrival of the end of the preamble (that is, for $1 \leq i \leq M - 1$), however, the scaling factor has a negative bias as an estimator for the system using a random preamble sequence [3]. Furthermore, the severity of the bias increases with an increase in i for $1 \leq i \leq M - 1$.

It follows that as the signal-to-noise ratio is increased, the adaptive-threshold technique only partially compensates for the stochastic increase in the test statistics. Thus it satisfies its design objective only imperfectly. Furthermore, the scaling factor is susceptible to the effect of random noise. The choice of the window size represents a tradeoff between limiting the effect of random noise on the scaling factor and its responsiveness to the change in the received input power upon the arrival of the preamble. It is shown in [3] that there is a finite optimal window size for any given set of system parameters.

The same considerations that motivated the introduction of the min-max performance criterion in Section 4.2 for fixed-threshold acquisition are also relevant for adaptive-threshold acquisition. In each system we consider using the adaptive-threshold technique, the nominal threshold η is chosen to achieve the min-max optimal probability of not acquiring for $\text{SNR} \geq \text{SNR}_{\text{min}}$ (where $\text{SNR}_{\text{min}} = 16$ dB in each example). In each example concerning adaptive-threshold acquisition, a window size of $W = 100$ is used.

We assume that for a packet arriving at time $t = 0$, the W test statistics necessary to form the scaling factor S_{M-Q} are available to the acquisition algorithm at time $t = (M-Q)T_c$. We also strengthen an assumption stated in Chapter 2 by assuming that there are no other transmissions present during the interval $[(M-Q-W)T_c, MT_c]$.

5.3 Performance with the Min-Max Optimal Nominal Threshold

The effectiveness of the adaptive-threshold technique is illustrated in Fig. 5.1, which corresponds to an example in [3]. The probability of not acquiring in a system with length-400 preamble sequences is shown for both fixed-threshold acquisition and adaptive-threshold acquisition. The use of the adaptive threshold markedly reduces the severity of the non-monotonicity of the performance as a function of the preamble signal-to-noise ratio. Moreover, the worst-case probability of not acquiring is decreased by more than an order of magnitude in comparison with fixed-threshold acquisition.

The choice of the window size provides a less desirable tradeoff if the preamble sequence is short than if it is long, and thus the adaptive threshold technique is less beneficial in a system with a short preamble. Suppose the system employs a preamble length of $M = 26$, for example. The optimal window size for the scaling factor is much larger than the preamble length, and the resulting adaptive-threshold algorithm yields almost the same worst-case performance as fixed-threshold acquisition as shown in Fig. 5.2.

The value of the adaptive-threshold technique for a system using intermediate-length preamble sequences falls somewhere between its value with short sequences and its value with long sequences. This is illustrated in Fig. 5.3 which shows the probability of not acquiring in a system with length-100 preamble sequences. The impact of the adaptive threshold on performance is not as pro-

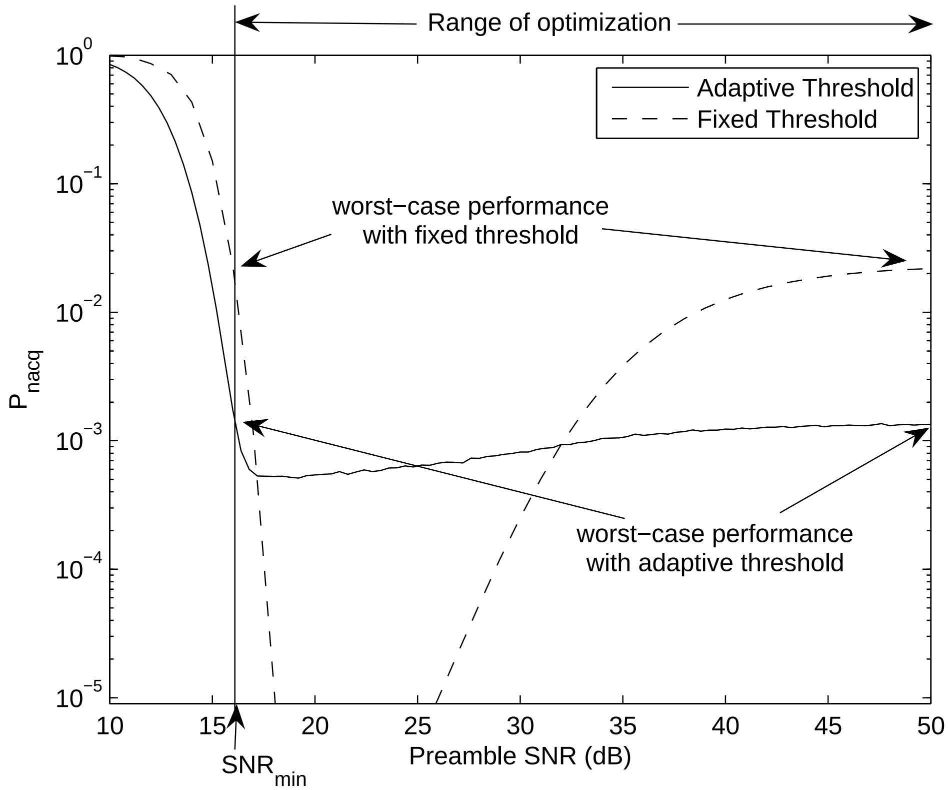

Figure 5.1: Comparison of fixed-threshold acquisition and adaptive-threshold acquisition for preamble sequences of length 400.

nounced as with the length-400 preamble sequences. However, the adaptive-threshold algorithm still provides a factor of 3.2 improvement in acquisition performance compared with fixed-threshold acquisition.

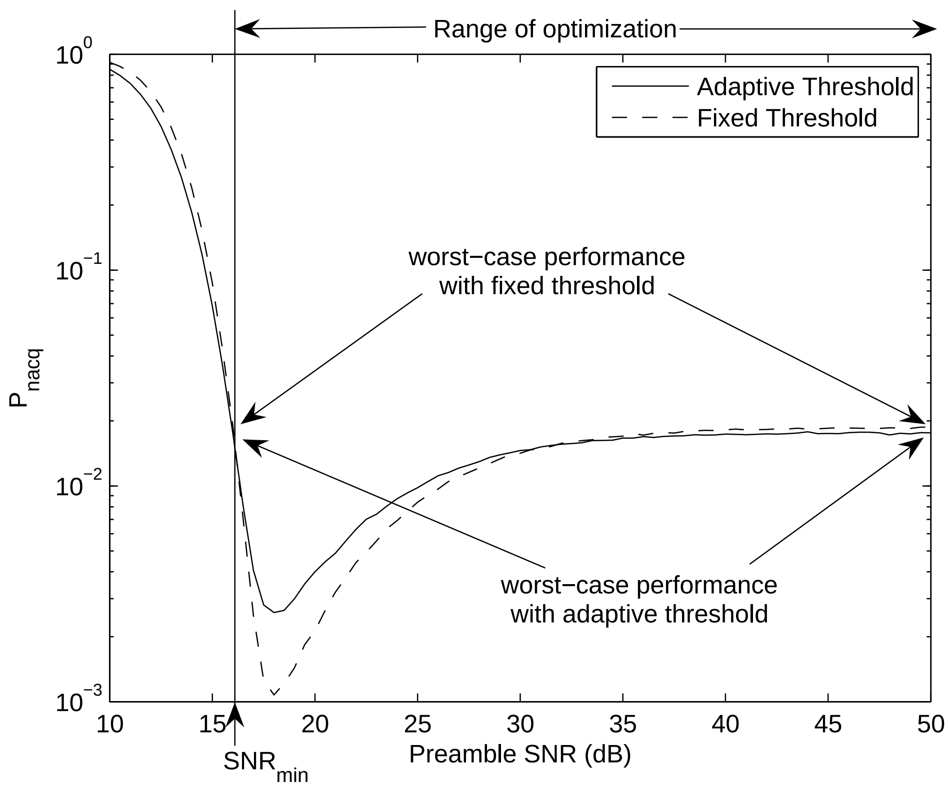

Figure 5.2: Comparison of fixed-threshold acquisition and adaptive-threshold acquisition for preamble sequences of length 26.

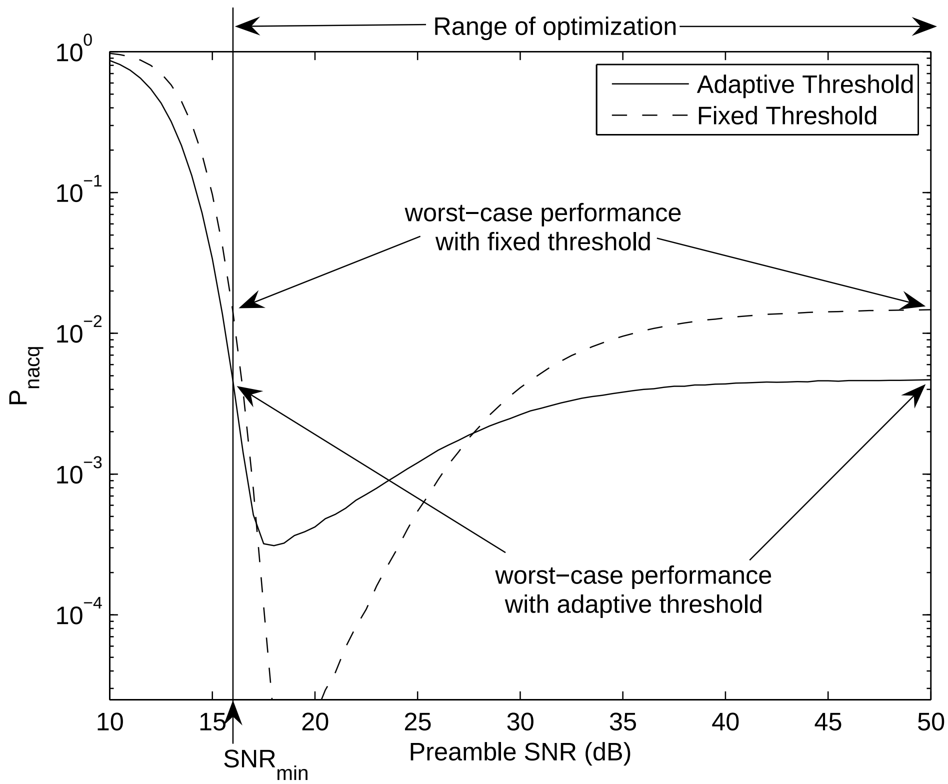

Figure 5.3: Comparison of fixed-threshold acquisition and adaptive-threshold acquisition for preamble sequences of length 100.

CHAPTER 6

SEQUENCE SELECTION FOR IMPROVED ACQUISITION PERFORMANCE

In this chapter, we present a simple metric that provides an approximate ranking of candidate preamble sequences with respect to the acquisition performance that results with their use. A system is considered in which only sequences satisfying an acceptance criterion based on the metric are used as preamble sequences. We consider the performance improvement in fixed-threshold acquisition that can be obtained for a given stringency in the acceptance criterion, and we examine the relationship between the preamble sequence length and the effectiveness of sequence selection based on the acceptance criterion.

6.1 Sequence-Selection Algorithm

The aperiodic autocorrelation function of the preamble sequence is an important factor in determining the acquisition performance of the system, as noted in Section 4.1. Thus it is desirable to identify a simple figure of merit based on the aperiodic autocorrelation function of a quaternary sequence that serves as a good predictor of the acquisition performance when the sequence is used as a preamble sequence. In this thesis, we focus on the *sidelobe energy* of the sequence [13] as its figure of merit. The sidelobe energy Z of a quaternary sequence $\{a_i\}$ of length M is given by

$$Z = \sum_{k=1}^{M-1} |C_a(k)|^2,$$

where $C_a(k)$ is given in (3.2).

The usefulness of the figure of merit is characterized by its relationship with the probability of false alarm for a given fixed acquisition threshold and a given value of the preamble signal-to-noise ratio. This relationship is illustrated in Fig. 6.1 for three choices of the acquisition threshold in the system with preamble sequences of length 26. For each threshold, the average probability of false alarm is shown as a function of the sidelobe energy of the preamble sequence for a high value of the preamble signal-to-noise ratio. That is, the probability of false alarm is averaged over the sequences with a given sidelobe energy.

There is a marked dependence of the average performance on the sidelobe energy, and the relationship is at least approximately monotonic for a given threshold. Of course some individual sequences with a higher sidelobe energy result in better performance than some other individual sequences with a lower sidelobe energy, so that the sidelobe-energy selection criterion does not provide direct control over the *worst-case* performance among the individual selected sequences. But it serves as an effective means to control the *average* performance over the set of selected sequences.

The sequence-selection algorithm considered in this thesis employs a pseudo-random sequence generator (which is modeled in the evaluations here as a

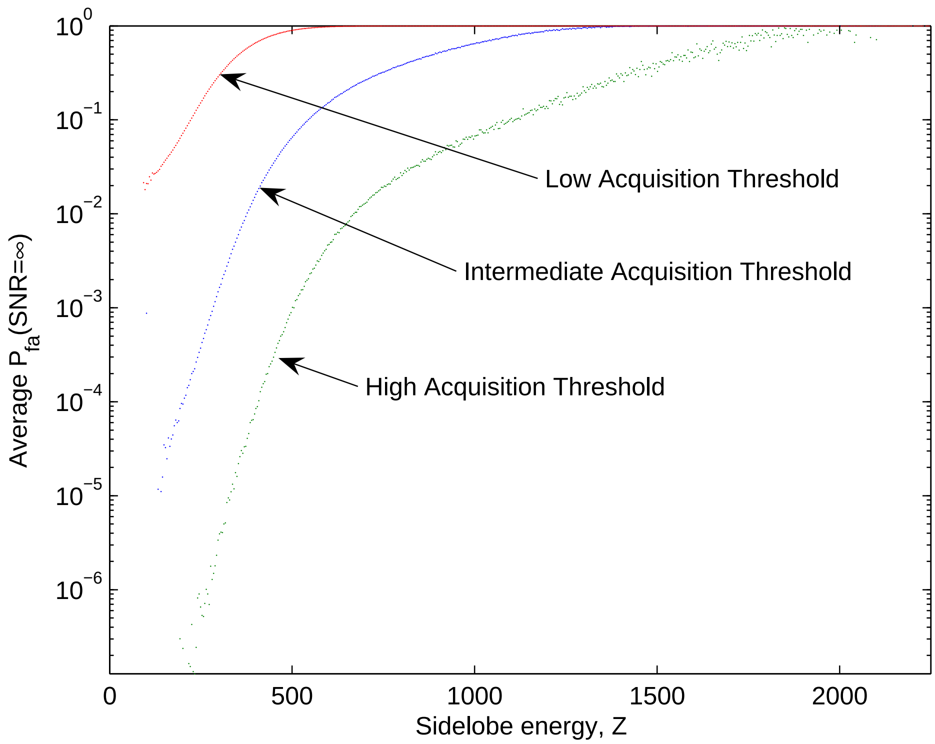

Figure 6.1: Average probability of false alarm as a function of the sidelobe energy Z.

random-sequence generator). Each generated sequence is tested for acceptance or rejection as the preamble sequence for the upcoming time epoch by comparing the sidelobe energy for the generated sequence with an *acceptance threshold* $Z_{\max}$. If the sidelobe energy is less than or equal to the acceptance threshold, the sequence is accepted. If the sidelobe energy is greater than the acceptance threshold, however, the sequence is rejected. If the sequence is rejected, the next candidate preamble sequence is generated according to the generation algorithm and the same test is applied to it. The process continues until a sequence is generated that passes the acceptance criterion.

All nodes generating the preamble sequence for the upcoming epoch use the same generation algorithm, the same initial state of the algorithm based on the identifier of the upcoming epoch, and the same acceptance criterion. Thus they all determine the same accepted sequence for use in the epoch. Because of the monotonic relationship between the sidelobe energy and the average acquisition performance, the choice of the acceptance threshold can be determined from a specified requirement for the average acquisition performance. Thus there is a well-defined tradeoff between the average acquisition performance over the accepted preamble sequences and the number of sequences accepted with a given generation algorithm.

6.2 Performance of Fixed-Threshold Acquisition with Sequence Selection

In this section, we consider the performance of serial, matched-filter acquisition using a fixed acquisition threshold and a preamble sequence that is determined by random-sequence generation followed by application of the sequence-acceptance criterion. Performance is given in terms of the average probability of not acquiring, where the average is over all preamble sequences that are generated using a specified acceptance threshold. For each choice of the acceptance threshold, the acquisition threshold is chosen to yield min-max optimal performance with $\text{SNR}_{\text{min}} = 16$ dB.

The sequence-selection algorithm can be used to substantially decrease the min-max probability of not acquiring for the system with short preamble sequences. This is illustrated in Fig. 6.2, which shows the acquisition performance for preamble sequences of length 26 and each of six values of the acceptance threshold. The values of the acceptance threshold are chosen such that the percentage of randomly generated quaternary sequences that satisfy the acceptance criterion ranges between 1% (i.e., 99% of the sequences are rejected) and 100% (i.e., none of the sequences are rejected). The figure also shows the min-max optimal performance for the ideal length-26 preamble sequence that has an (unrealizable) sidelobe energy of zero.

The use of all randomly generated preamble sequences (i.e., the absence of any sequence selection) results in a worst-case probability of not acquiring of

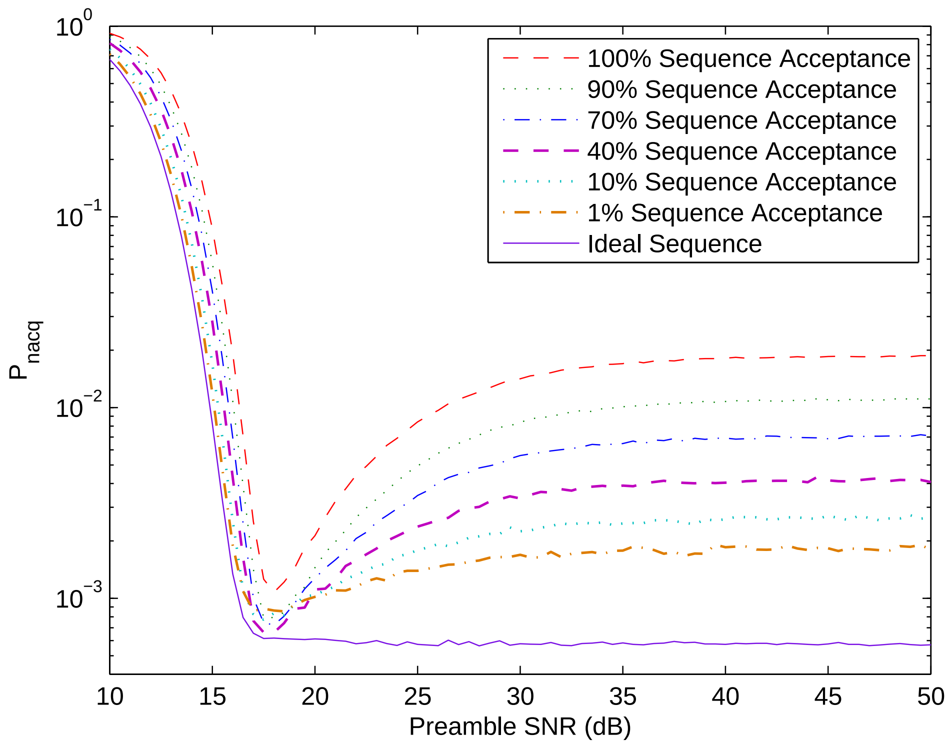

Figure 6.2: Performance with length-26 preamble sequences and fixed-threshold acquisition for various acceptance percentages.

0.02. If the preamble sequences are restricted to the 1% of length-26 sequences with a sidelobe energy in the lowest 1% among all such sequences, the worst-case probability of not acquiring is only 0.002. Thus exclusion of 99% of the possible preamble sequences results in an order-of-magnitude reduction in the worst-case probability of not acquiring. The performance obtained with the 1% acceptance criterion is close to the worst-case probability of error with the ideal sequence, which is 0.0134. Thus there is only limited opportunity for further improvement by application of a more stringent acceptance criterion to the length-26 preamble sequences.

A similar improvement in performance is obtained if sequence selection is used in the system with intermediate-length preamble sequences. This is shown in Fig. 6.3 for length-100 preamble sequences and the same six acceptance percentages considered in Fig. 6.2. The performance with the ideal sequence of length 100 is also shown. A reduction in the acceptance percentage from 100% to 1% results in reduction in the min-max probability of not acquiring from 0.01462 to 0.00163. Once again, the error probability is decreased by almost a factor of ten. The worst-case probability of not acquiring is only 0.00054 with the ideal length-100 sequence, however. Thus there may be a somewhat better opportunity for improvement by use of an acceptance criterion of less than 1% with the length-100 preamble sequences than was true with the length-26 preamble sequences.

For long preamble sequences, the sequence-selection algorithm yields a much more modest improvement in the min-max probability of not acquiring for a given stringency in the acceptance criterion. This is illustrated in Fig. 6.4 for preamble sequences of length 400. For the system using short preamble sequences, a 1% sequence acceptance rate results in a min-max probability of not acquiring that is decreased by a factor of 10.4 compared with the performance if a 100% acceptance rate is used. In contrast, reducing the acceptance rate from 100% to 1% decreases the min-max error probability by a factor of only 2.5 if length-400 preamble sequences are used. There is also a large difference between the worst-case probability of error with the 1% acceptance criterion

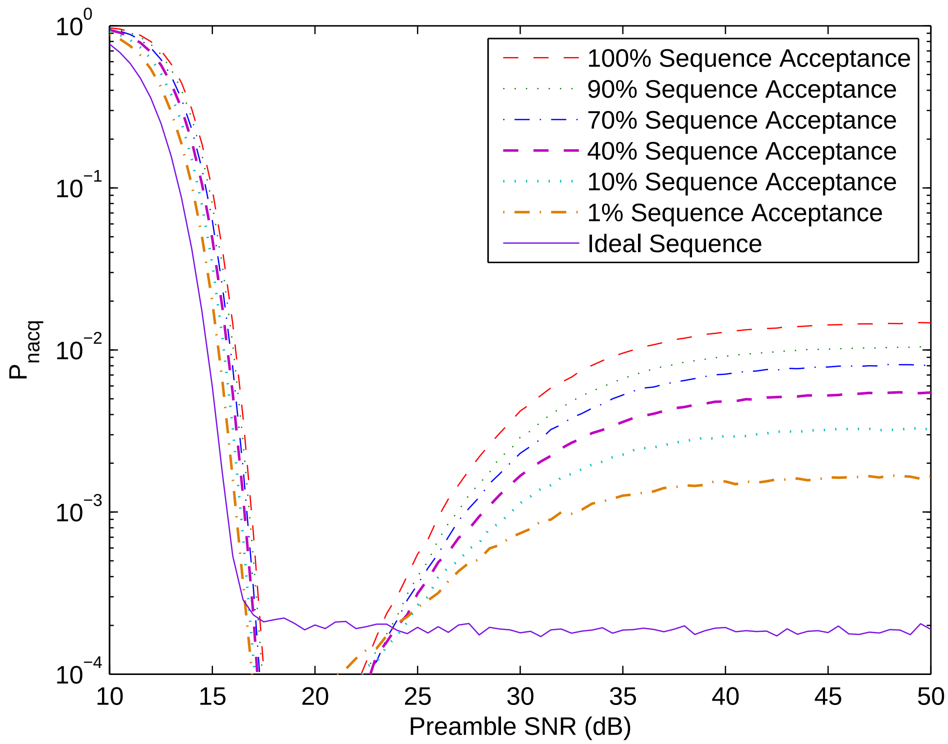

Figure 6.3: Performance with length-100 preamble sequences and fixed-threshold acquisition for various acceptance percentages.

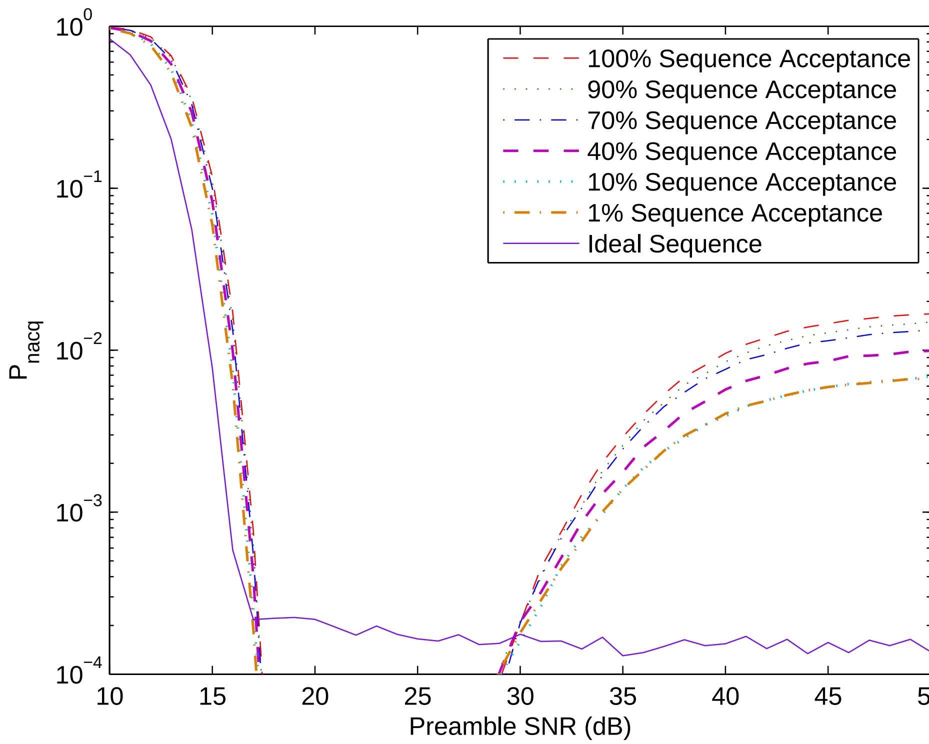

Figure 6.4: Performance with length-400 preamble sequences and fixed-threshold acquisition for various acceptance percentages.

and the worst-case probability of error with the ideal length-400 preamble sequence. The two values are 7×10^{-3} and 5.8×10^{-4}, respectively.

It is clear that as the sequence acceptance criterion is made more stringent, a larger number of candidate preamble sequences must be evaluated on average to find an acceptable sequence and the computation required to find an acceptable sequence increases. In some applications it may be important to minimize the computation required for this task. The computational burden imposed by the sequence-selection algorithm in any given epoch is determined by the number of candidate sequences that must be tested in order to find an

Acceptance Percentage	**worst-case $\mathbf{P_{nacq}}$**	**Search Run Length**	
		Mean	**Std. Dev.**
90%	0.0109	1.11	0.345
80%	0.0086	1.23	0.533
70%	0.0068	1.41	0.748
60%	0.0058	1.65	1.02
50%	0.0049	2.00	1.40
40%	0.0043	2.49	1.91
30%	0.0036	3.34	2.78
20%	0.0031	4.90	4.37
10%	0.0026	9.20	8.72
1%	0.0019	99.8	98.1

Table 6.1: Performance and run-length statistics for various acceptance percentages.

acceptable sequence, which is referred to as the *search run length*. Clearly, the minimum search run length is one.

The worst-case probability of not acquiring and the mean and standard deviation of the search run length resulting from random sequence generation of short preambles are shown in Table 6.1 for various acceptance percentages. It is apparent that there is a significant trade-off between the acquisition performance achieved with sequence selection and the computational burden imposed by the sequence-selection algorithm. Note that the mean and standard deviation of the search run length are in approximately inverse proportion to the acceptance percentage.

6.3 Performance of Adaptive-Threshold Acquisition with Sequence Selection

In this section, we consider the performance of a system using the sequence-selection algorithm in conjunction with the adaptive-threshold technique. Performance is given in terms of the average probability of not acquiring, where the average is over all preamble sequences that are generated using a specified acceptance threshold. As in the results of the previous section, the acquisition threshold is chosen to yield min-max optimal performance with $\text{SNR}_{\min} = 16$ dB for each choice of the acceptance threshold.

The acquisition performance of a short-preamble system using the adaptive-threshold technique is improved substantially by the incorporation of a stringent sequence-selection criterion. This is illustrated in Fig. 6.5 which shows the acquisition performance for the system using adaptive-threshold acquisition with preamble sequences of length 26 and each of six values of the acceptance threshold. The figure also shows the min-max optimal performance for the ideal length-26 preamble sequence that has a sidelobe energy of zero. The min-max probability of not acquiring is decreased from 0.01762 to 0.00368 as the acceptance percentage is reduced from 100% to 1%, a decrease of slightly less than five-fold.

For a system with intermediate-length preamble sequences, the sequence-selection algorithm provides a moderate improvement in acquisition performance when using an adaptive threshold. This is illustrated in Fig. 6.6 which

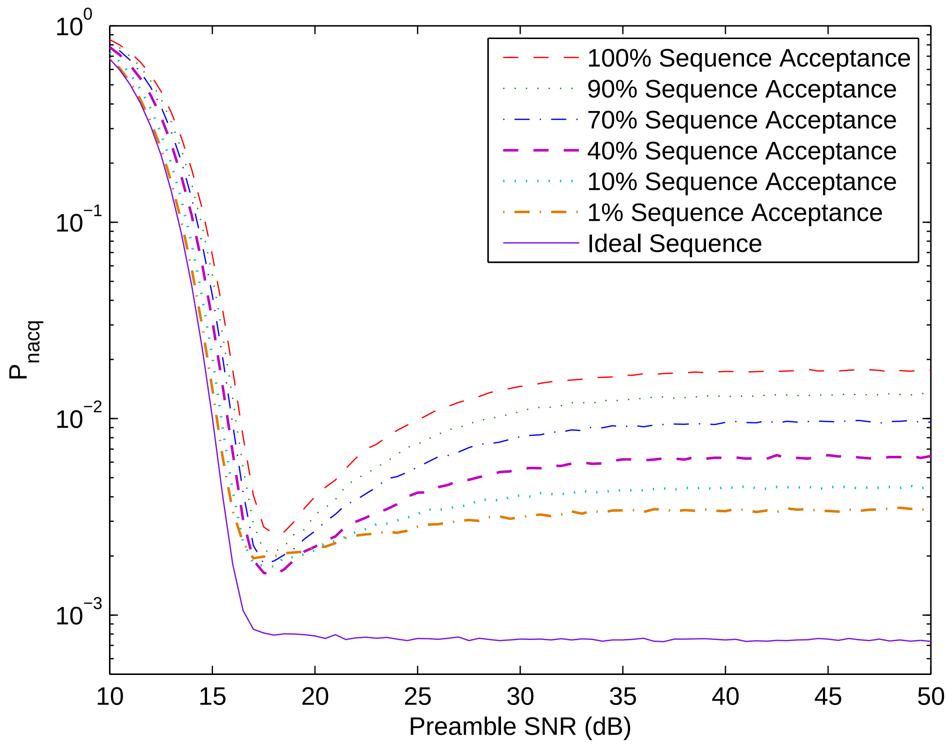

Figure 6.5: Performance with length-26 preamble sequences and adaptive-threshold acquisition for various acceptance percentages.

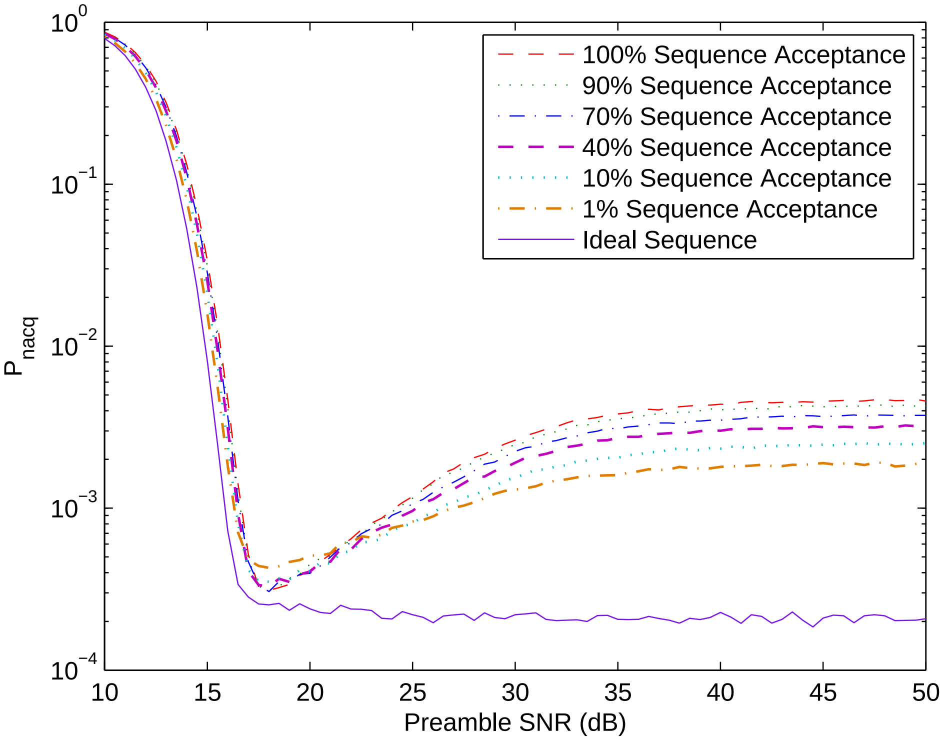

Figure 6.6: Performance with length-100 preamble sequences and adaptive-threshold acquisition for various acceptance percentages.

shows the acquisition performance for the system using adaptive-threshold acquisition with preamble sequences of length 100 and each of six values of the acceptance threshold. The figure also shows the min-max optimal performance for the ideal length-100 preamble sequence that has a sidelobe energy of zero. The min-max probability of not acquiring is decreased from 0.00464 to 0.00187 as the acceptance percentage is reduced from 100% to 1%, a decrease of 2.5-fold.

Utilization of a sequence-selection criterion has much less impact on the performance of the adaptive-threshold technique for the long-preamble system

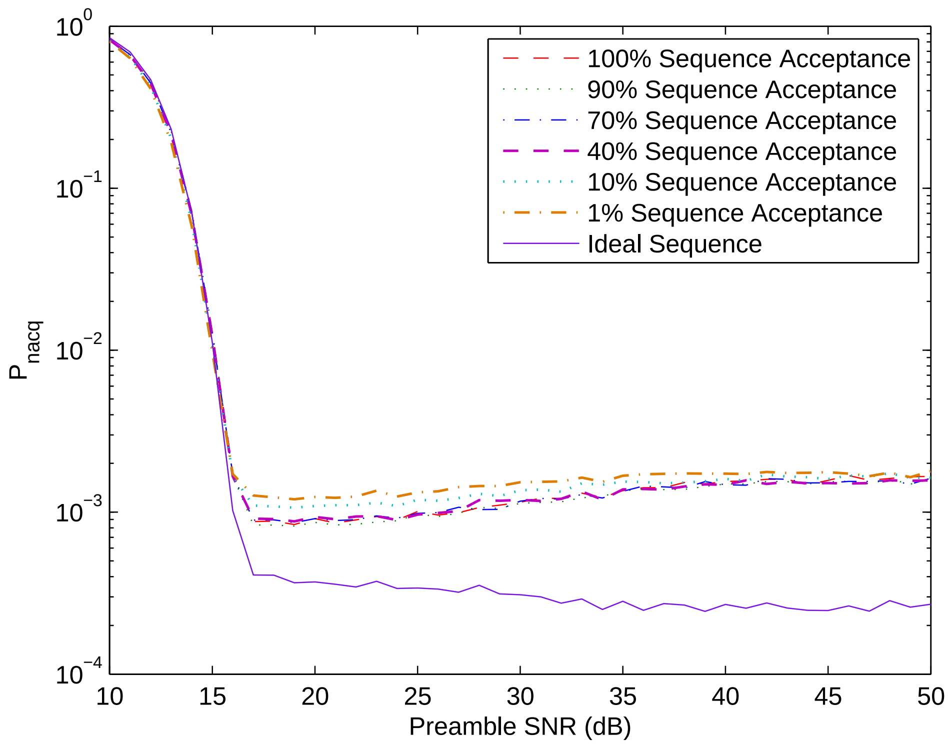

Figure 6.7: Performance with length-400 preamble sequences and adaptive-threshold acquisition for various acceptance percentages.

than for either the short and intermediate-length preamble systems. This is shown in Fig. 6.7 (which corresponds to a figure in [14]) for the system with preamble sequences of length 400 and adaptive-threshold acquisition. For acceptance percentages as low as 1%, sequence selection has a negligible effect on the worst-case probability of not acquiring.

6.4 Effectiveness of Sequence Selection

The usefulness of the sequence-selection technique depends heavily on sequence length because the distribution of the sidelobe energy differs for different sequence lengths. This is illustrated by considering $\tilde{Z}$, which is the *normalized sidelobe energy* given by

$$\tilde{Z} = \frac{Z}{\mathrm{E}[Z]}.$$

The probability mass function of the normalized sidelobe energy for each of the three preamble lengths considered in this thesis is shown in Fig. 6.8. The figure shows that as the preamble sequence length is increased, the distribution of the normalized sidelobe energy becomes more closely centered around the mean. Thus, as the preamble sequence length is increased, the impact of even a stringent sequence acceptance criterion has a decreased effect on the average sidelobe energy of the pool of accepted sequences. Consequently, the usefulness of the sequence-selection technique is diminished as the preamble sequence length is increased.

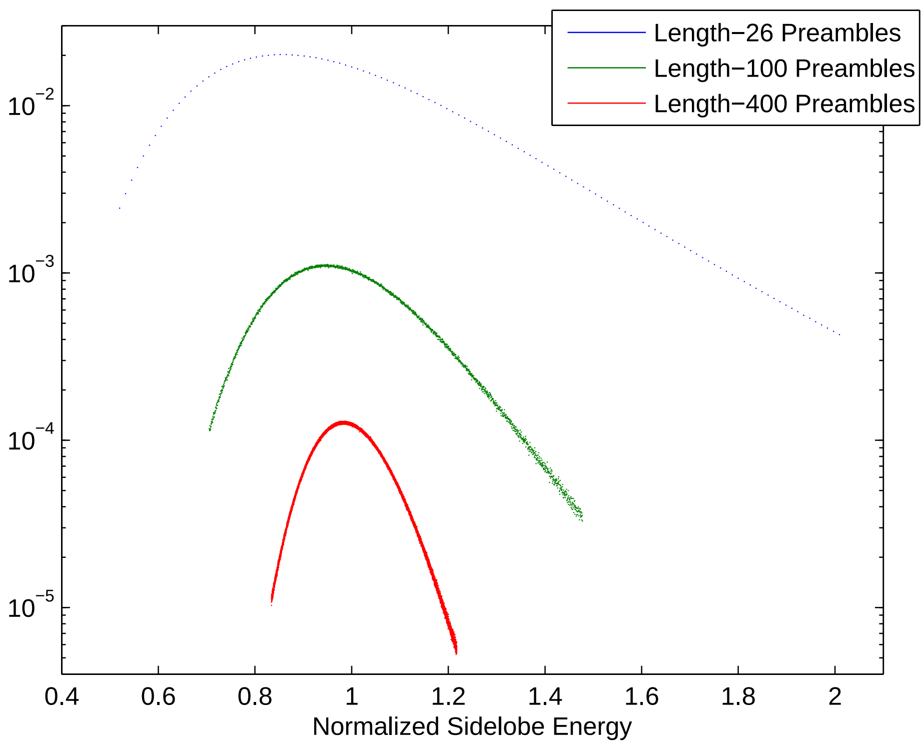

Figure 6.8: Probability mass functions of the normalized sidelobe energy for the three preamble sequence lengths considered.

CHAPTER 7

COMPARISON OF ACQUISITION TECHNIQUES

In this chapter, we compare the performance obtained using four approaches to serial, matched-filter acquisition based on the two techniques discussed in the previous chapters. The four approaches include the use of a sequence-selection criterion alone, the use of the adaptive-threshold technique alone, the use of the two in combination, and the use of neither. Performance is illustrated for a system with a short preamble sequence, a system with an intermediate-length preamble sequence, and a system with a long preamble sequence. The values of β, SNR_{min}, and W are the same as for the examples in the previous chapters. The combinations of preamble length and verification interval are also the same as in the previous chapters, except where noted in Section 7.4 with respect to the verification interval.

7.1 Performance with Short Preamble Sequences

The use of a stringent acceptance criterion for preamble sequences substantially improves the acquisition performance of a short-preamble system regardless of whether fixed-threshold acquisition or adaptive-threshold acquisition is employed, as shown in the previous chapter. The benefit of sequence selection is much less pronounced if the adaptive-threshold technique is used than if the fixed-threshold technique is used. Furthermore, the use of the adaptive-threshold technique provides only a negligible benefit if all preamble sequences are accepted, as shown in Chapter 5. Consequently, the use of an adaptive threshold is actually detrimental to performance if a stringent sequence-acceptance criterion is imposed. This is illustrated in Fig. 7.1 which shows the acquisition performance for both fixed-threshold acquisition and adaptive-threshold acquisition with sequence acceptance percentages of 100% and 1%. The performance of fixed-threshold acquisition with an ideal preamble sequence is also shown.

The acquisition performance of the short-preamble system with the adaptive-threshold technique is marginally better than with the optimal fixed threshold if all randomly generated preamble sequences are accepted, and the worst-case probability of not acquiring is approximately 0.02 for both. The worst-case probability of not acquiring is decreased by a factor of 10.4 as the acceptance percentage is reduced from 100% to 1% in the system using a fixed acquisition threshold, however, whereas a corresponding decrease of a factor of only 5.2 occurs for the system using an adaptive acquisition threshold. Thus, fixed-threshold acquisition results in a worst-case probability of not acquiring of 0.0018 if 1% sequence-acceptance criterion is used, whereas the corresponding probability of not acquiring is 0.0034 for adaptive-threshold acquisition.

The performance of the system using an adaptive threshold depends on the window size used in the adaptive-threshold algorithm. Thus the relative performance of the two acquisition techniques for a given acceptance percent-

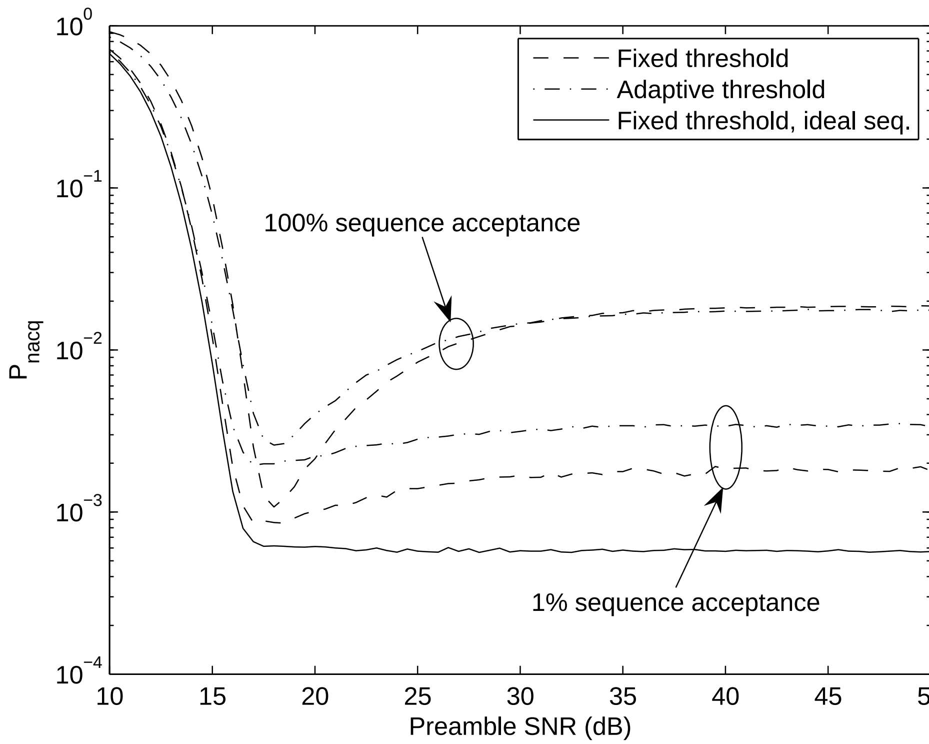

Figure 7.1: Performance with length-26 preamble sequences for four combinations of acquisition technique and acceptance percentage.

age depends on the same parameter. For the system with short preamble sequences and an adaptive threshold, the tradeoff implicit in the choice of the window size favors a window size that is large relative to the length of the preamble sequence. Thus with 100% sequence acceptance, little improvement in performance can be obtained by reducing the window size to below 100.

The sidelobe energy of the accepted sequences has less impact on the acquisition performance as the stringency of the selection criterion is increased, and the tradeoff represented by the choice of the window size favors an even larger size for the adaptive-threshold technique. Consequently, better acquisition performance is achieved in the short-preamble system for a stringent acceptance criterion if a very large window size is employed. But in that case, the performance is almost identical to the performance obtained with the fixed acquisition threshold. Thus for the combination of verification interval and short preamble in this example, the adaptive-threshold technique is of little value and the optimal fixed acquisition threshold results in nearly the best performance regardless of the acceptance percentage.

7.2 Performance with Intermediate-Length Preamble Sequences

The use of a stringent acceptance criterion improves the acquisition performance substantially for a system with intermediate-length preambles and fixed-threshold acquisition, but the improvement provided by sequence selection is more modest if the system uses adaptive-threshold acquisition. Yet the

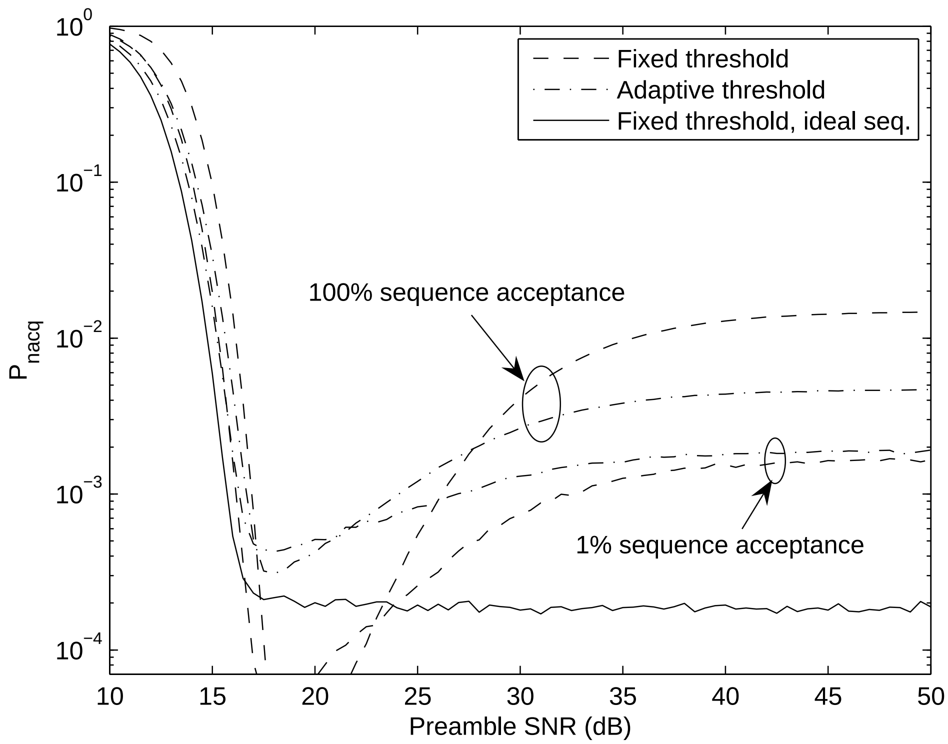

Figure 7.2: Performance with length-100 preamble sequences for four combinations of acquisition technique and acceptance percentage.

adaptive-threshold technique provides a substantial improvement over fixed-threshold acquisition if all preamble sequences are accepted. The net result of these factors is similar worst-case acquisition performance for the systems using fixed-threshold acquisition and adaptive-threshold acquisition if a stringent sequence-acceptance criterion is imposed. This is illustrated in Fig. 7.2.

The acquisition performance of the intermediate-length-preamble system is much better with the adaptive-threshold technique than with the fixed-threshold technique if all sequences are accepted. The worst-case probability of not acquiring is 0.015 for the latter but only 0.005 for the former, a difference of a factor of three. The worst-case probability of not acquiring is decreased by a factor of 8.95 as the acceptance percentage is reduced from 100% to 1% in the system using a fixed acquisition threshold, but the corresponding decrease is only 2.49 for the system using an adaptive acquisition threshold. Thus, fixed-threshold acquisition results in a worst-case probability of not acquiring of 0.0018 if 1% sequence-acceptance criterion is used, which is slightly better than the probability of not acquiring for adaptive-threshold acquisition with the same acceptance percentage.

A comparison of the performance of the system using a fixed threshold and the system using an adaptive threshold favors the former more heavily as the acceptance percentage is decreased. An acceptance percentage between 10% and 100% results in better performance with adaptive-threshold acquisition, but an acceptance percentage below 1% results in better performance with fixed-threshold acquisition. Somewhat better performance can be obtained with adaptive-threshold acquisition if the window size is optimized for each acceptance percentage. For the same reasons addressed in the previous section, however, fixed-threshold acquisition will still yield better performance with a sufficiently stringent acceptance criterion.

7.3 Performance with Long Preamble Sequences

The use of a stringent acceptance criterion provides only a modest improvement in the performance of a system with length-400 preambles and

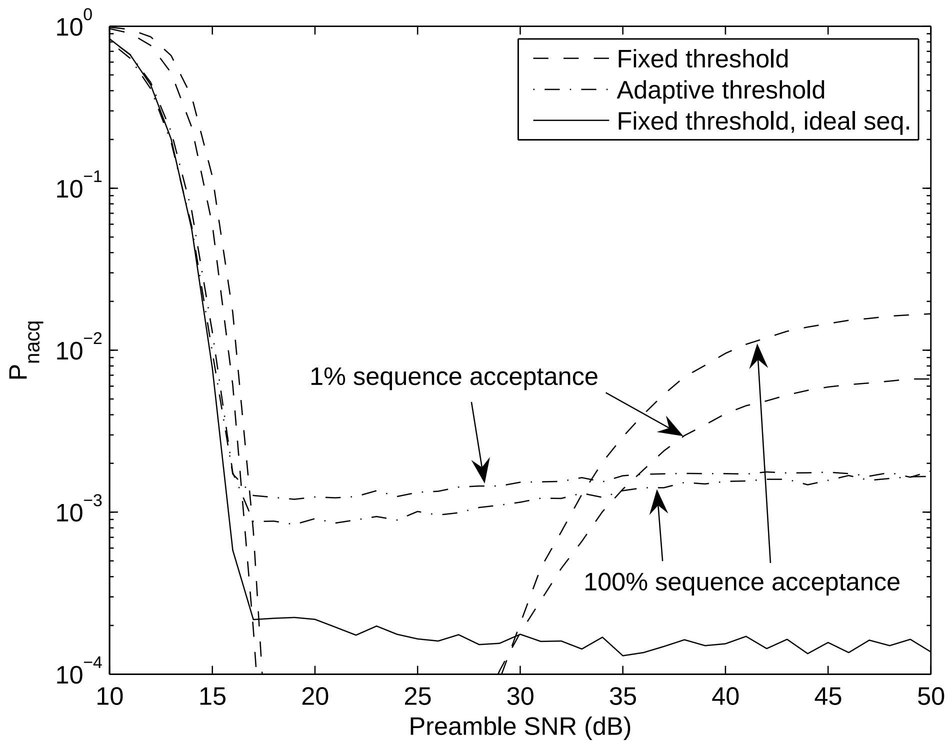

Figure 7.3: Performance with length-400 preamble sequences for four combinations of acquisition technique and acceptance percentage.

fixed-threshold acquisition, and it results in a negligible improvement if the system uses adaptive-threshold acquisition. The adaptive-threshold technique provides a dramatic improvement over fixed-threshold acquisition if all preamble sequences are accepted. Consequently, adaptive-threshold acquisition results in better performance than fixed-threshold acquisition over a wide range of sequence acceptance percentages, though its benefit diminishes as the acceptance criterion is made more stringent. This is illustrated in Fig. 7.3.

As observed in Chapter 5, the acquisition performance of the long-preamble system with the adaptive-threshold technique is an order of magnitude better than its performance with the fixed-threshold technique if all sequences are accepted. The worst-case probability of not acquiring is decreased by a factor of two as the acceptance percentage is reduced from 100% to 1% in the system using a fixed acquisition threshold, but the corresponding decrease is negligible for the system using an adaptive acquisition threshold. Thus, adaptive-threshold acquisition results in a worst-case probability of not acquiring of 0.0018 if 1% sequence-acceptance criterion is used, which is about one-fourth of the probability of not acquiring for fixed-threshold acquisition with the same acceptance percentage. Thus for the combination of verification interval and long preamble in this example, sequence selection of moderate stringency is of essentially no value and the best performance is achieved with an adaptive threshold alone.

7.4 Effect of Verification Interval on Performance

The examples in the previous chapters and sections employ a verification interval QT_c for which $Q = 2.5M$ for each system. Numerical results are also considered for two values of the verification interval: one with $Q = M$, and the other with $Q = 5M$. The different verification intervals are considered with short preamble sequences and intermediate-length preamble sequences. All other parameters are as in the earlier examples, except that for each choice of the verification interval, the corresponding min-max optimal threshold is used.

Over the range of verification intervals considered, the worst-case probability of not acquiring exhibits little sensitivity to the verification interval. The sensitivity is somewhat greater in the system with short preamble sequences than in the system with preamble sequences of intermediate length. It is also somewhat greater with a more stringent sequence-acceptance criterion than with a less stringent one. It is also greater with the adaptive-threshold technique than with the fixed-threshold technique. Even for the system with short preamble sequences, an adaptive threshold, and a 1% acceptance criterion, however, the worst-case probability of error increases by only a factor of two as Q is increased from 26 to 130. Thus the conclusions drawn in previous chapters and sections are applicable over a wide range of values of the verification interval.

CHAPTER 8

CONCLUSIONS

Two methods to improve the acquisition performance of a packet radio system are investigated. It is shown that adaptive-threshold acquisition and preamble-sequence selection both improve the performance of serial, matched-filter acquisition in packet radio communications, though different combinations of the two techniques are beneficial under different circumstances. Acquisition performance is evaluated for short, intermediate-length, and long preamble sequences.

If the system parameters correspond to those typical of narrowband packet-radio communications (i.e., short preamble sequences), only preamble-sequence selection is of practical value in mitigating the effect of the IF filter and AGC system on acquisition performance. The use of a moderately selective sequence-acceptance criterion can result in a decrease in the min-max probability of not acquiring by more than an order of magnitude compared with the use of all randomly generated preamble sequences.

If instead the parameters correspond to those typical of DS spread-spectrum packet-radio communications (i.e., long preamble sequences), the use of an adaptive acquisition threshold is much more valuable than preamble-sequence selection for the levels of selection stringency considered in this thesis. The use of an adaptive acquisition threshold alone results in performance nearly as good as that obtained using the two techniques in combination in that instance. Adaptive-threshold acquisition can result in a decrease in the min-max probability of not acquiring by more than an order of magnitude compared with optimal fixed-threshold acquisition.

If the packet radio system employs intermediate-length preamble sequences, the best acquisition performance is observed for a system employing both adaptive-threshold acquisition and preamble-sequence selection for most values of the sequence acceptance percentage. The individual contribution to the improvement in acquisition performance due to the adaptive-threshold technique is not as great as for the system with long preamble sequences. Similarly, the individual contribution to the improvement in acquisition performance due to preamble-sequence selection is less significant than for the short-preamble system. The combination, however, results in a decrease in the min-max probability of not acquiring of almost an order of magnitude compared with optimal fixed-threshold acquisition.

REFERENCES

[1] U. Madhow and M. B. Pursley, “Acquisition in direct-sequence spread-spectrum communication networks: An asymptotic analysis,” *IEEE Trans. Inform. Theory*, vol. 39, pp. 903–912, May 1993.

[2] D. L. Noneaker, A. R. Raghavan, and C. W. Baum, “The effect of automatic gain control on serial, matched-filter acquisition in direct-sequence packet radio communications,” *IEEE Trans. Veh. Technol.*, vol. 50, no. 4, pp. 1140–1150, July 2001.

[3] A. Swaminathan and D. L. Noneaker, “A technique to improve the performance of serial, matched-filter acquisition in direct-sequence spread-spectrum packet radio communications,” *IEEE J. Select. Areas Commun.*, vol. 23, no. 5, pp. 909–919, May 2005.

[4] J. Schloemann and D. L. Noneaker, “Performance of serial, matched-filter packet acquisition using a preamble-sequence acceptance criterion,” in *Proc. IEEE Military Commun. Conf.*, Washington, DC, Oct. 2006.

[5] M. B. Pursley, “The role of spread spectrum in packet radio networks,” in *Proc. IEEE*, vol. 75, no. 1, Jan. 1987, pp. 116–134.

[6] D. V. Sarwate and M. B. Pursley, “Crosscorrelation properties of pseudorandom and related sequences,” *Proc. IEEE*, vol. 68, no. 5, pp. 593–619, May 1980.

[7] K. K. Chawla and D. V. Sarwate, “Acquisition of PN sequences in chip synchronous DS/SS systems using a random sequence model and the SPRT,” *IEEE Trans. Commun.*, vol. 42, no. 6, pp. 2325–2334, June 1994.

[8] B. Razavi, *RF Microelectronics*. Upper Saddle River, New Jersey: Prentice Hall, 1998.

[9] M. E. Van Valkenburg, *Analog Filter Design*. New York: Oxford University, 1982.

[10] D. L. Noneaker, “The performance of serial, matched-filter acquisition in direct-sequence packet radio communications,” in *Proc. IEEE Military Commun. Conf.*, McLean, VA, Oct. 2001, pp. 1045–1049.

[11] G. L. Turin, “The characteristic function of Hermitian quadratic forms in complex normal variables,” *Biometrika*, vol. 47, pp. 199–201, June 1960.

[12] M. Schwartz, W. R. Bennett, and S. Stein, *Communication Systems and Techniques*. New York: McGraw-Hill, 1966.

[13] M. B. Pursley and H. F. A. Roefs, “Numerical evaluation of correlation parameters for optimal phases of binary shift-register sequences,” *IEEE Trans. Commun.*, vol. 27, no. 10, pp. 1597–1604, Oct. 1979.

[14] J. Schloemann and D. L. Noneaker, “Preamble sequence selection for adaptive-threshold, serial, matched-filter acquisition in packet radio communications,” in *Proc. 2007 IEEE Internat. Symp. Commun. Theory and Applications*, Ambleside, UK, July 2007.