Digital Signal Processing 51 (2016) 1–9

Contents lists available at ScienceDirect

Digital Signal Processing

www.elsevier.com/locate/dsp

The constant false alarm rate property in transformed noncoherent detection processes

G.V. Weinberg

National Security, Intelligence, Surveillance and Reconnaissance Division, Defence Science and Technology Group, P. O. Box 1500, Edinburgh, South Australia 5111, Australia

ARTICLE INFO

Article history:
Available online 22 January 2016

Keywords:
Constant false alarm rate
Radar detection
Pareto distributed clutter
Transformations

ABSTRACT

Recently a transformation approach for noncoherent radar detector design has been introduced, where the classical constant false alarm rate detectors for Exponentially distributed clutter are modified to operate in any clutter intensity model of interest. Recent applications of this approach have introduced new decision rules for target detection in Pareto and Weibull distributed clutter. These transformed detectors tended to lose the constant false alarm rate property with respect to one of the clutter parameters. A closer examination of this transformation process yields conditions under which the constant false alarm rate property can be retained. Based upon this, a new model for X-band maritime radar returns is investigated, and corresponding detectors are developed. The relative merits of this new development are investigated with synthetic and real X-band data.

1. Introduction

1.1. Constant false alarm rate detectors

Constant false alarm rate (CFAR) detectors are of considerable importance in maritime surveillance radar signal processing, and their development continues to be explored in many different clutter environments [1–8]. Such detectors are usually proposed without formulation of the optimal Neyman–Pearson detector, but are often derived based upon analysis of clutter distribution properties and range-Doppler map characteristics [9,10]. Due to the fact that Neyman–Pearson detectors require knowledge of clutter distribution parameters and target strength, these alternative detection processes are often introduced in an attempt to avoid variation in the desired false alarm probability by sacrificing detection performance [11,12].

In earlier low-resolution radar systems, the assumption of Exponentially distributed amplitude squared (or intensity) clutter was found to be valid, and as a result of this, it was possible to apply a simple detection process to introduce CFAR. In particular, based upon a range-Doppler map, measurements either in range or Doppler (or both) can be taken such that they are sufficiently separated from a cell under test (CUT), and then averaged. This average is then normalised using a threshold multiplier, which provides adaptive control of the false alarm probability. The presence of a target in the CUT is declared if the return being tested exceeds this normalised measure [11].

Fig. 1 illustrates this process in more general terms, and is implemented as a sliding window which runs across the range-Doppler map [13]. The detection strategy takes a sample of observations (in this case in range or Doppler) and passes them after square law detection (SLD) to a shift register as shown. The returns are separated into a series of components. Measurements of the clutter are subdivided into two disjoint sets $\{C_1, C_2, \ldots, C_M\}$ and $\{C_{M+1}, \ldots, C_{M+N}\}$, each of length M and N respectively. These are separated from the CUT by a number of guard cells (shown as shaded in Fig. 1). Based upon the two subsets of clutter, two measurements are taken, denoted by f_1 and f_2, which provide a synthesis of the clutter, via an appropriate statistic (such as averaging or order statistics, for example). These two measurements are then combined to produce an overall measurement of clutter, denoted by $g(f_1, f_2)$. The latter is normalised by τ, which is used to regulate the false alarm probability. This is then compared with the CUT and a detection decision can be made. This binary result is recorded, and the sliding window is then continued over the range-Doppler map, and the results can be applied to a moving target indicator to allow tracking [14].

In the case where the clutter measurements are homogeneous and independent Exponentially distributed returns, a detection process such as that illustrated in Fig. 1, can achieve the CFAR property for a very large class of admissible clutter measurement

E-mail address: Graham.Weinberg@defence.gov.au.

http://dx.doi.org/10.1016/j.dsp.2016.01.005
1051-2004/

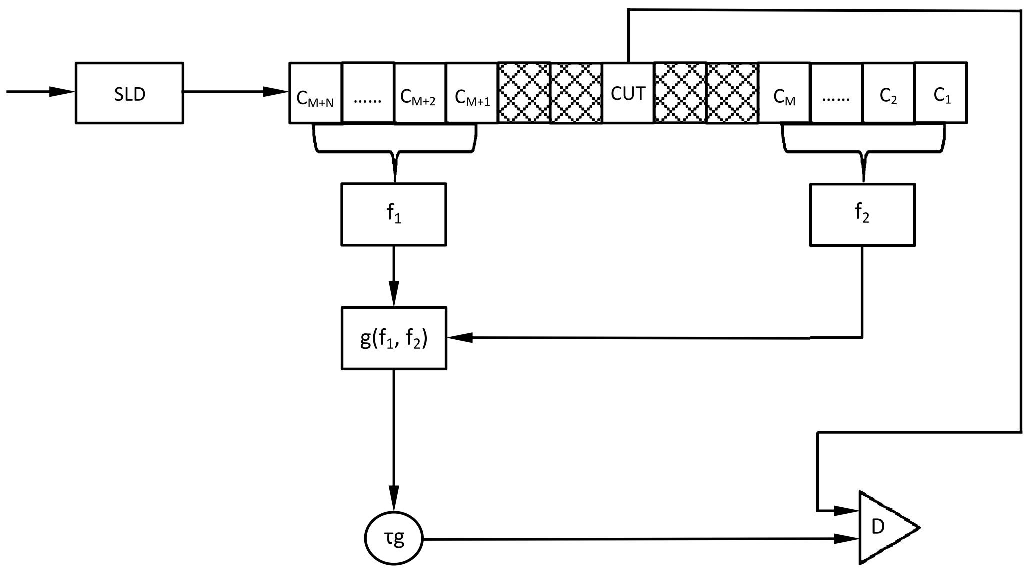

Fig. 1. A sliding window detection scheme which is run across a range-Doppler map to arrive at a binary detection decision. The complex signal is processed with square law detector (SLD), and then passed to a shift register as shown. Clutter cells (denoted by $C_1, C_2, \ldots, C_{M+N}$) are separated from the cell under test (CUT) using guard cells as shown. The clutter statistics are then processed, normalised and compared to the CUT to arrive at a binary detection decision. These decisions then provide input to tracking algorithms and moving target indicators.

functions f_1, f_2 and $g(f_1, f_2)$ [15]. The selection of alternatives to the averaging case have been proposed due to the fact that as the sliding window is run across the range-Doppler map, the data used for clutter measurement will be subjected to interference and changes in the clutter homogeneity. Interference can be due to secondary targets present in the clutter cells, which can be managed through the selection of an order statistic (OS) for f_1 and f_2, for example. In the context of maritime surveillance radar, changes in the clutter homogeneity can arise from variations in the sea surface. Hence, although a simpler alternative to the Neyman–Pearson based detectors, it is an engineering challenge to design an effective sliding window detection scheme that achieves CFAR and manages these performance issues. This is a much more difficult exercise when such processes are constructed for higher resolution maritime surveillance radar.

1.2. Development of CFAR

The cell-averaging CFAR (CA-CFAR) [16], which bases the measurement of clutter level on averaging, is optimal for target detection with a Swerling 1 target model in homogeneous Exponentially distributed clutter [17]. However, in the presence of clutter irregularities and interference, the CA-CFAR experiences severe performance degradation [15]. Hence there have been many alternative CFAR schemes proposed. Earlier attempts to improve on the CA-CFAR's performance involved the idea of taking the smallest-of (SO) of two cell-averaged components of the range profile [18]. This detector, known as the SO-CFAR, arose out of the need to improve the resolution capability of the CA-CFAR. The greatest-of (GO) CFAR was introduced in an attempt to better manage false alarm regulation [11,19]. However, both these approaches have severe drawbacks as outlined in [15]. The idea of applying an order statistic to measure the average level of clutter was thus introduced in [13], producing an OS-CFAR, which could be designed to manage interfering targets very well, while also regulating the false alarm probability. A further improvement on the OS-CFAR is introduced in [15], where a trimmed mean is used. This involves censoring of a number of lower and upper clutter measurements in the range profile, and then using the sum of the remainder as the estimate of the clutter power level. More recent advances in CFAR for target detection in Exponentially distributed clutter can be found in [20], where a switching mechanism is used to improve performance of the CA-CFAR.

With improvements in radar resolution, there has been a natural progression from the Exponential intensity clutter model to others that capture the statistical structure of higher resolution radar's clutter returns. Consequently, there has been extension of the CFAR schemes designed for Exponential intensity clutter to the Log-Normal [21], Weibull [3,22–24], K [5,25,26] and Pareto [27,28,30] distributed clutter environments. In the extension to such clutter models, two general approaches have been used. The first is to apply the CFAR scheme developed for Exponentially distributed returns directly to the new model of interest. This then necessitates the determination of the appropriate threshold multiplier for the new clutter environment [24,27]. The second approach is to exploit characteristics of the clutter model [21,22]. The first approach can always be applied, while the second depends on the availability of useful properties of the underlying clutter model.

1.3. Contributions and structure

This paper is concerned with a recently developed transformation approach for CFAR detector development, introduced in [32], which was developed in an attempt to provide a systematic way in which detectors could be produced. This approach begins with the well-developed CFAR detectors designed to operate in independent and identically distributed Exponential clutter, and introduces a mapping allowing such detectors to be transformed to operate in any clutter model of interest. This generalised the results of [28], which examines the special case of the transformation of detectors to operate in Pareto distributed clutter. Also illustrated in [32] is the result of transforming CFAR detection processes to operate in Weibull distributed clutter. The advantages of this transformation approach is that detectors can be very easily produced for new clutter environments while retaining the original formulation of the threshold multiplier. As shown in [27], direct adaptation of classical detection schemes to the Pareto case can add computational complexity in the threshold multiplier calculation. The disadvantage of the transformation method is that clutter Parameter dependence may occur, resulting in the loss of the CFAR property with respect to a particular parameter.