Generic Ground Clutter Simulation for Radar Testing and Evaluation

Prepared by:

Johannes Jurgens Strydom

STRJOH039

Prepared for:

Prof. M. Inggs

Radar Remote Sensing Group

Department of Electrical Engineering
University of Cape Town

External study leader:

J. E. Cilliers

Council for Scientific and Industrial Research

17 August 2012

Submitted to the Department of Electrical Engineering at the University of Cape Town in partial fulfilment of the academic requirements for a

Master of Engineering *specialising in Radar*

DECLARATION

I know the meaning of plagiarism and declare that all the work in the document, save for that which is properly acknowledged, is my own.

Signature of Author: _____

Johannes Jurgens Strydom

17 August 2012

ABSTRACT

Hardware in the loop (HIL) simulation is often used to reduce the cost of evaluating the performance of a radar [1]. These hardware in the loop simulators aim to recreate an accurate representation of the environment of a radar. Traditionally radar environment simulators focused on single scatterer targets, and electronic countermeasures, but as the fidelity of the radar increases these simulation methods alone become less useful. Thus the fidelity of radar environment simulators has to be increased if they are to remain relevant. Targets modelled using multiple complex scatterers have been incorporated into HIL simulators, and radar clutter simulation is the next step in this process of increasing fidelity, which is the focus of this study.

The design challenge is to balance the contradictory requirements of complexity and fidelity of a clutter simulation system. The fidelity has to be high enough to accurately simulate cutter with realistic properties so as to be indistinguishable from true clutter returns for the radar under test, but the complexity has to be low enough to execute the clutter simulation on a field programmable gate array (FPGA) platform in real time. Before the FPGA simulation can be performed, the algorithms have to be designed and tested in software. This project aims to investigate the real time approximation of relevant distributions that should be used when testing and evaluating a radar in clutter scenarios. Simulation of these relevant distributions are investigated to determine their accuracy.

First a literature survey is presented to analyse the aspects of clutter that should be considered for incorporation into a radar environment. The most commonly used statistical models for surface clutter were chosen and methods for their simulation from literature were identified that could potentially be implemented on an FPGA. The actual FPGA implementation did not form part of this project, but was a constraint during the evaluation of the methods. The methods identified were implemented in software for analysis. The probability density function (PDF), spectral shape, autocorrelation function, and higher order moments were analysed. Furthermore, the trends pertaining to the effects of correlation in the samples, and the effects of varying the most important distribution parameters were analysed. A new approach to simulate the statistics of radar clutter was then developed to answer the question of whether the methods identified in literature could be improved upon. The new approach was tested and compared to the methods in literature, and a path for possible future research in this regard emerged.

The most frequently used distributions for clutter modelling identified were, the Rayleigh, Log-Normal and Weibull distributions. These distributions can be simulated in different ways. A critical part of simulating the distributions for a radar environment is the addition of pulse to pulse correlation. Memoryless non-linear transforms (MNLT) were chosen because they transform a correlated Gaussian distribution to the required distribution while approximately maintaining the input's correlation characteristics. Spherically invariant random vector (SIRV) methods were excluded from analysis because they produce samples in batches, and are generally mathematically more complex. Both the complexity, and batch generation property, makes them less attractive for the future goal of implementation on FPGA based systems. The Rayleigh distribution did not require any transform, since the I and Q components are Gaussian distributed. The Rayleigh distributed samples generated in this way performed well, and served as a benchmark for the other methods, because no transform from the Gaussian distributed values was required. The method chosen to generate the correlated Log-Normal distributed values performed well, but will be relatively complex to implement in hardware. The method chosen for the Weibull distributed sample generation exhibited inaccuracies. The Weibull PDF shape degraded beyond a useful state for certain settings of correlation.

The poor performance for the Weibull method, and the complexity of the Log-Normal method prompted the question of what the performance would be for the simplest possible solution to simulating correlated clutter? Since the methods from literature were insufficient, a different approach was developed. This method showed similar results for all three chosen distributions. The problems experienced with the Weibull distributions were not present with this new method, and the complexity is much lower than that of the Log-Normal method. A draw-back was that the transformation to the required distribution decorrelated the samples relative to the input correlation.

The new approach was found to be of low enough complexity to allow for future development for implementation on an FPGA, and showed advantages in this regard, relative to the methods from literature. The problem for decorrelation could be addressed with pre-calculations and correction before it is eventually implemented on an FPGA, but requires further research. Open questions have been identified and the areas where possible future research into this method is required are highlighted.

ACKNOWLEDGEMENTS

I would like to thank my supervisor, Professor Mike Inggs and the University of Cape Town for the opportunity to complete my masters degree through the new M.Eng degree programme. I would also like to thank Professor Mike Inggs for his support and guidance, from the small administrative processes, to assistance with the courses, and all the way through to the completion of this dissertation.

I would like to thank my mentor Jacques Cilliers for his support which I have come to rely on heavily, throughout the past four years. Your contribution to my professional and personal development is much appreciated.

I would like to thank my family and friends for their support. Trying to balance studies, work and my social responsibilities was challenging. Without your assistance and understanding it would not have been possible to come all this way, with you still by my side.

I would like to thank my employer, the CSIR, for its assistance financially, and with study time throughout the year. I would like the thank DPSS and the Radar and Electronic Warfare group for the opportunity to align my masters degree with my projects, which contributed tremendously to this effort.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF SYMBOLS

A - Magnitude
b - K-Distribution scale parameter
BW - Bandwidth
c - Speed of light
e - Mathematical constant (Euler's number)
E_k - Expected value from histogram bin number k
f_d - Doppler frequency
f_o - Center frequency
Fs - Sampling frequency
g - Gaussian random variable
g_n - Gaussian random sample number n
G - Antenna gain
h - Height
$H(f)$ - Transfer function
H_0 - Null hypotheses
H_1 - Alternate hypotheses
I - Complex signal In-phase channel
j - Imaginary unit
k - Weibull distribution shape parameter
k_T - Total number of bins in a histogram
K_v - Modified Bessel function
m - Metres
n - Moment order
n - Sample number
N - Total number of samples
O_k - Observed value from histogram bin number k
p - Number of parameters for distribution under test for χ^2 test
P_a - Analytical PDF
P_{ac} - AC power
P_m - Measured PDF
P_r - Power received
P_t - Power transmitted
Q - Complex signal Quadrature channel

r - Rate parameter

r - Ratio between AC and DC power

R - Range

$S(f)$ - Frequency Spectrum

S - Surface area

t_p - Pulse length

v - Degrees of freedom for χ^2 test

v - K-Distribution shape parameter

v - Velocity

w - Wind speed

x - Instantaneous amplitude of probability density function

$\overline{x}$ - Sample mean

x_n - Sample number n generated with the correlation filter

x_{n-1} - Previous sample generated with the correlation filter

$X(z)$ - Z-transform filter input variable

$Y(z)$ - Z-transform filter output value

$^\circ$ - Angle in degrees

α - Alpha stable distribution characteristic exponent

α - Rayleigh distribution scale parameter

α_c - Critical angle

α_d - Depression angle

α_g - Grazing angle

α_i - Incidence angle

β - Beta parameter for the correlation filter

β - Exponential AC Shape Parameter

χ^2 - Chi-Square test

δ - Delta Dirac function

η - Eta parameter for the correlation filter

γ - Alpha stable distribution dispersion parameter

γ - Surface type constant for constant gamma model

Γ - Gamma function

λ - Weibull distribution scale parameter

λ - Wavelength

μ - Mean

ϕ - Phase of complex signal

π - Mathematical constant (Pi)

δ - Alpha stable distribution location parameter

σ - Radar cross section

σ - Standard deviation

σ_h - RMS surface height deviation from the average

σ^0 - Clutter radar cross section

$\tau_{1/2}$ - Half decorrelation time

$\tau_{1/e}$ - $1/e$ decorrelation time

θ_H - Horizontal beam width

θ_V - Vertical beam with

NOMENCLATURE

AC	-	Alternating Current
dB	-	Decibel
CDF	-	Cumulative Density Function
CI	-	Confidence Interval
DC	-	Direct Current
DFT	-	Discrete Fourier Transform
DSP	-	Digital Signal Processing
FFT	-	Fast Fourier Transform
FPGA	-	Field Programmable Gate Array
HIL	-	Hardware In the Loop
ICDF	-	Inverse Cumulative Distribution Function
IID	-	Independent and Identically Distributed
KL	-	Kullback-Liebler
MIT	-	Massachusetts Institute of Technology
MNLT	-	Memoryless Non-Linear Transform
NA	-	Not Applicable
NRCS	-	Normalised Radar Cross Section
PDF	-	Probability Density Function
PRF	-	Pulse Repetition Frequency
radar	-	RAdio Detection And Ranging
RCS	-	Radar Cross Section
RES	-	Radar Environment Simulator
RMS	-	Root Mean Square
RSP	-	Radar Signal Processor
RV	-	Random Variable
SIRV	-	Spherically Invariant Random Vector
ZMNL	-	Zero Memory Non-Linearity

1 INTRODUCTION

A radar transmits electro-magnetic energy into an environment, and receives the returning echoes to identify objects in this environment. Not all echoes are necessarily of interest to the radar. Radar clutter is interference from objects in the environment that are not of interest with respect to the function of the specific radar. The performance of radars designed to detect aircraft is diminished by echoes from the ground, sea, rain, mist, clouds, birds, and depending on the sensitivity of the radar, even insects. Clutter might also be intentionally introduced into the environment by non cooperative targets deploying chaff[1]. Clutter can be divided into three categories: point clutter, surface clutter and volume clutter. Surface clutter can be further divided into ground clutter and sea clutter. The problem addressed in this project was the simulation of realistic ground clutter for radar testing and evaluation.

Testing new radar algorithms in the physical world is a time consuming task due to the continuously changing nature of the environment. It is also extremely difficult or even impossible to artificially re-create the occurrence of a rare environmental event. This non-repeatability induces a degree of uncertainty in the validity of results. A reliable simulation method for producing clutter is thus required to test the radar in a repeatable signal environment. Some realism is necessarily sacrificed to simulate the environment in real time, and for the ability to manipulate and control this simulated environment. Control over the simulation will allow engineers to enhance and improve the effectiveness of their radar algorithms with repeated and standardized testing in a simulated environment that includes realistic correlated ground clutter.

Hardware in the loop (HIL) simulations can be used to reduce the development time and development cost of a radar system. Hardware in the loop simulation based testing can be performed before testing in the field. This allows the radar designer to test the radar signal processor (RSP) before the optimization of weight, processing and power for the integration into a larger more complex system occurs. The radar environment simulator (RES) system developed at the CSIR[2] can perform these types of tests by simulating targets, and electronic counter measures. Clutter simulation is the next step in extending the functionality of this system, improving the radar

1 A cloud of electromagnetic reflective material ejected and spread in the surrounding air by an aircraft as a countermeasure to confuse a radar.

2 The Council for Scientific and Industrial Research (CSIR) in South Africa is one of the leading scientific and technology research, development and implementation organisations in Africa (www.csir.co.za).

testing and evaluation capability of the HIL system. Research is required to investigate and develop effective clutter generation algorithms for the current architecture, which is based on field programmable gate arrays (FPGA).

The technical challenge of this project was to obtain and develop methods which can simulate the stochastic nature of clutter in real time, while adhering to the constraints of the limited resources available on the hardware platform. The implementation of the design will be for eventual integration into existing firmware written for an FPGA based hardware platform. The underlying principles for FPGA based processors have to be accommodated into the design of the clutter simulation algorithms. Special attention is thus needed to identify the operations that map easily to an FPGA architecture, and similarly to identify the ones that do not. Operations that do not map easily to an FPGA will have to be modified, approximated or completely re-designed in order for the algorithm to perform efficiently on the FPGA.

While complying with the constraints of the FPGA architecture the algorithms have to produce results that are statistically accurate while at the same time be easily controllable. The probability density function (PDF) has to be accurate with regards to shape while still maintaining the specified spectral characteristics. The autocorrelation function also has to have the shape described in literature. Eventual implementation on an FPGA with fixed point mathematical operations will not have the same dynamic range as that of floating point operations on a general purpose computer. An additional limitation is that computers use 32 or even 64 bits to represent numbers, while the FPGA uses as few as possible to save logic resources. Some FPGA devices contain digital signal processing (DSP) blocks, which are mathematical operations that are implemented as optimised single function digital logic gates on the chip, as opposed to reconfigurable logic. These resources also have limitations on the maximum number of bits with which calculations can be executed.

This study on clutter simulation and clutter simulation design is limited to the identification and implementation of different clutter distributions and their underlying behaviour. The effects of hardware implementation (such as limited dynamic range and resolution) on the resulting distributions are also investigated. The simulation scenario is limited to a stationary ground based radar, receiving reflections from ground based surfaces, such as trees and soil. The radar is assumed to have a single and constant look direction, and effects such as radar scanning are not included in this simulation. Simulation of internal errors due to the design of radar, such as sidelobes from the radiation pattern of the antenna, are also not incorporated into the objectives of this project.

1.1 Objectives

This project aims to identify and develop algorithms for clutter simulation which are suitable for implementation on an FPGA based hardware platform. The clutter algorithms have to generate clutter according to the statistics used to model clutter, which have been identified by means of a literature survey. The autocorrelation function, and thus frequency spectrum of clutter, has to be simulated while meeting the requirements for the shape of the modelled PDF. A relatively basic frequency spectrum is assumed, to reduce complexity. The frequency spectrum must have a 3 decibel (dB) bandwidth in line with a desired value set for the -3 dB bandwidth before the simulation starts. The simulation methods have to be coded and tested for eventual implementation on an FPGA based hardware platform. Thus special attention should be paid to algorithm complexity with respect to the resources that are available on current FPGAs.

The objectives of the project are thus to simulate

- pulse to pulse correlated ground clutter, with an
- accurate PDF, and an
- accurate spectral -3 dB bandwidth.

The implementation of the algorithms on an FPGA platform does not form part of the scope of this project, due to the time constrains imposed on the project. However, attention was given to the practicalities of implementing the algorithms for the eventual simulation on an FPGA platform.

1.2 Executive summary

First a literature survey on the topic of ground clutter is presented in Section 2. The literature survey covers all the important aspects that are required to describe ground clutter. The topics cover radar cross section (RCS), spatial distribution, amplitude distribution, spectral distribution, correlation properties, wavelength and polarization properties of clutter. For clutter simulation the RCS, Doppler spectrum and amplitude distribution have been identified as the most important aspects to consider, since the remaining parameters can be translated to a scaling of the final clutter power. From all these aspects described a basic foundation for clutter simulation can be formed. The Rayleigh, Log-Normal and Weibull distributions have been identified as the most widely used models for ground clutter.

Two approaches were followed. The first approach was the evaluation of complex accurate algorithms from literature, followed by the second approach which was the development and testing of a new algorithm of lower complexity.

Section 3 covers the implementation of radar clutter simulation, which aims to replicate a realistic radar clutter environment. The task of the simulation is to generate data that represents theoretical radar clutter, while still being stochastic in nature. Most computers available today cannot generate true randomness. For this reason methods have been developed to create pseudo-randomness. Computer programs such as Matlab have built in uncorrelated random number generators, which are high fidelity and stochastically well behaved. These random number generators can be used for stochastic simulation on a computer, but are not necessarily compatible with hardware implementation. Section 3.1 gives a short overview of the random number generation field, identifying four random number generation methods. These methods are the CDF inversion method, transformation method, rejection method, and the recursive method. Section 3.2 presents a derivation for the generation of correlated Gaussian samples. The resulting equations allows the designer to set the Doppler bandwidth of the generated clutter samples. The resulting first order filter of equation (1) (equation (30) on page 27) is used as the basis for all the clutter simulations in this project.

$$x_n = \eta x_{n-1} + \beta g_n \tag{1}$$

Section 3.3 describes the approach followed to evaluate the clutter generators that are presented in the sections that follow. The frequency spectrum, autocorrelation function and amplitude distribution were evaluated for each of the clutter generators. The χ^2 test and KL divergence value have been used in conjunction with the calculation of the higher order moments to analyse how well a generated result matches that of the analytical model.

The methods to generate correlated clutter by means of a memoryless non-linear transform (MNLT[3]) have correlated Gaussian random number generators as their input. The built in Matlab uncorrelated Gaussian random number generator is transformed into a correlated Gaussian random number generator, and tested in Section 3.4. These test results serve as a basis against which the performance of more complicated distributions are tested. The χ^2 test and KL divergence value identified a decrease in the PDF accuracy with an increase in correlation as shown in Figures 1 and 2 receptively (Figure 37 on page 42 and Figure 39 on page 43).

The Rayleigh distribution follows from a Gaussian distribution without any modification and was tested in Section 3.5. MNLT methods are used to generate Log-Normal and Weibull distributed clutter in sections 3.6 and 3.7. For each of the clutter distributions the PDF is measured, and compared to the analytical PDF. The -3 dB bandwidth of the random number sequence is measured

3 A memoryless non-linear transform (MNLT) is also referred to as a zero memory non-linearity (ZMNL) in some literature. For this document the acronym MNLT is used to refer to this type of transformation.

in the frequency domain. The autocorrelation function is used to obtain the time (measured in samples) it takes the signal to decorrelate by half, and to decorrelate by $1/e$. The behaviour of the performance of the PDF shape and accuracy of the -3 dB bandwidth is also investigated for variation in bandwidth and the specific distribution's shape parameters.

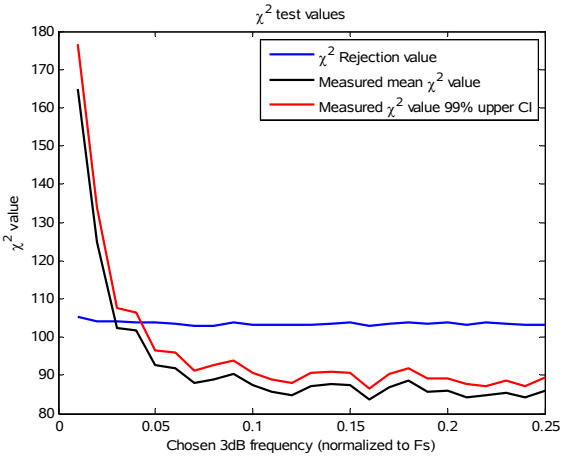

Figure 1: Plot of the outcome from a chi squared test.

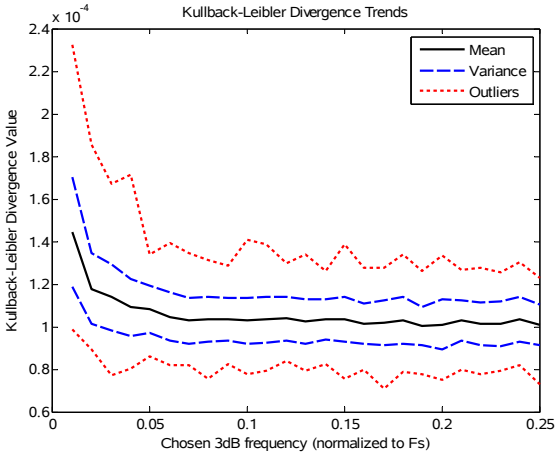

Figure 2: Kullback-Leibler divergence value.

The χ^2 test and KL divergence value for all three these distributions show the same trend of increased inaccuracy with a smaller bandwidth, similar to the Gaussian case of Section 3.4. The performance of the Rayleigh distribution was better than the Gaussian. The Log-Normal distribution was on the edge of rejection by the χ^2 test and the Weibull distribution was rejected most of the time. The KL divergence values did not show similar performance changes, and was very similar for all the distributions. The -3 dB bandwidth values were accurate across the entire range of tested values, and some deviation was noticeable at small bandwidth settings. Varying the shape parameters of the distributions resulted in poorer performance for settings where the shape parameter resulted in a longer tailed distribution. These methods from literature performed well, but will be difficult to implement on an FPGA due to the large number of mathematical operations.

A clutter simulation approach is also developed in Section 4 by investigating the operation for some of the signal processing steps performed by the typical radar signal processor. A new cumulative density function (CDF) transformation method is developed for the generation of clutter distributions in this section. This approach is specifically developed with the complexity of implementation on an FPGA hardware platform in mind. Section 4.1 describes the generation of the uniformly distributed phase, which is used in each of the experiments in the sections which follow. The Rayleigh distribution (Section 4.2), Log-Normal distribution (Section 4.3) and Weibull distribution (Section 4.4) are simulated using this approach, and tested using the same methods as used for the generation techniques from literature. The CDF transformation method was less

accurate than the methods from literature in terms of the PDF shape. The bandwidth of the generated clutter samples deviated strongly from the desired values as can be seen in Figure 3 (Figure 180 from page 113). Fortunately the deviation from the ideal was monotonic, which allows for its use regardless, since a single input value results in a single output value. The CDF transformation method however exhibited advantages such as lower complexity and greater reconfigurability (the same structure can be used for any distribution), which would make it more suitable for eventual FPGA implementation.

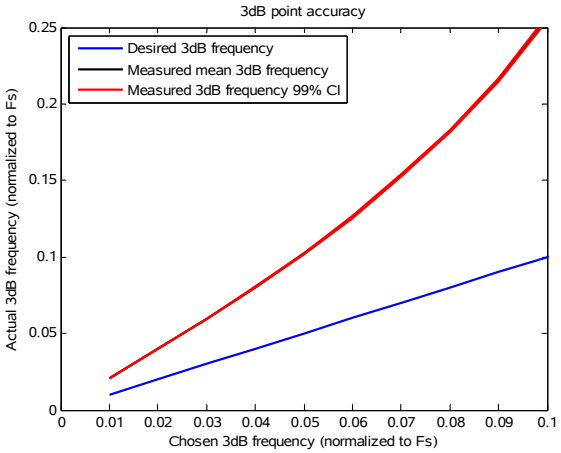

Figure 3: Measured -3 dB bandwidth of the correlated Rayleigh random number.

Finally the methods found in literature are compared to those of the new CDF transformation method in Section 5. A discussion of the measurement methods, in addition to future work is presented in the conclusion of this work as part of the same section.

2 LITERATURE SURVEY

The various aspects of radar clutter have been the subject of many studies and measurement campaigns. Barton [2] lists a number of parameters that must be specified when clutter is included in the performance evaluation of a radar. These parameters are:

- clutter radar cross section (RCS),
- number of discrete scattering sources,
- spatial extent and distribution of sources,
- velocity extent and distribution (Doppler spectrum),
- wavelength dependence of RCS,
- amplitude distribution (probability density function),
- spatial correlation of amplitudes, and
- polarization properties.

The impact of these aspects have to be investigated before clutter can be simulated realistically. Before any investigation can start the relationship between the different functional parts that influence the performance of the radar system is required. The radar equation (2) describes the power received P_R by a radar in relation to the power transmitted P_T, the gain of the antenna G, the wavelength of the transmitted signal λ, the radar cross section of the object σ, and the range from the radar R.

$$P_R = \frac{P_T G^2 \lambda^2 \sigma}{(4\pi)^3 R^4} \tag{2}$$

The ground clutter radar cross section is an extension of the radar equation for the reflectivity from ground surfaces.

2.1 Ground Clutter Radar Cross Section

Radar cross section is denoted by the symbol σ. It is the measure used to describe the magnitude of the reflectivity of an object. The unit of measurement for radar cross section is area in metres squared (m^2). This is the perceived size of the object by the radar, and not the physical size of the object. The perceived size is a strong function of the geometry, and reflective properties of the object. Larger objects that are less reflective will have a smaller radar cross section than a small object that is highly reflective. When the reflective object is a clutter source, the variable σ^0 is used to describe the magnitude of the reflectivity. For ground clutter the radar cross section is

dimensionless because the reflectivity of the ground is described as reflectivity per ground surface area, both measured in m^2. For clarity the literature uses m^2/m^2 This is known as the normalised radar cross section (NRCS). The surface area of the ground, within a radar resolution cell, is captured with the variable S as shown in (3). The product of the normalised radar cross section and the surface area provides the radar cross section of a ground surface.

$$\sigma = \sigma^0 S \tag{3}$$

The mean NRCS of clutter is characterised over incidence angles of 0° to 90° and can be split into three different regions. Figure 4 shows a plot for typical mean NRCS values between incidence angles of 0° and 90°. Incidence angles between 80° and 90° are classified as the grazing or near grazing region, angles between 30° and 80° are classified as the plateau region and angles between 0° and 30° are in the near vertical region [3]. For clutter analysis and modelling the mean of these statistics is not sufficient information to describe the clutter completely.

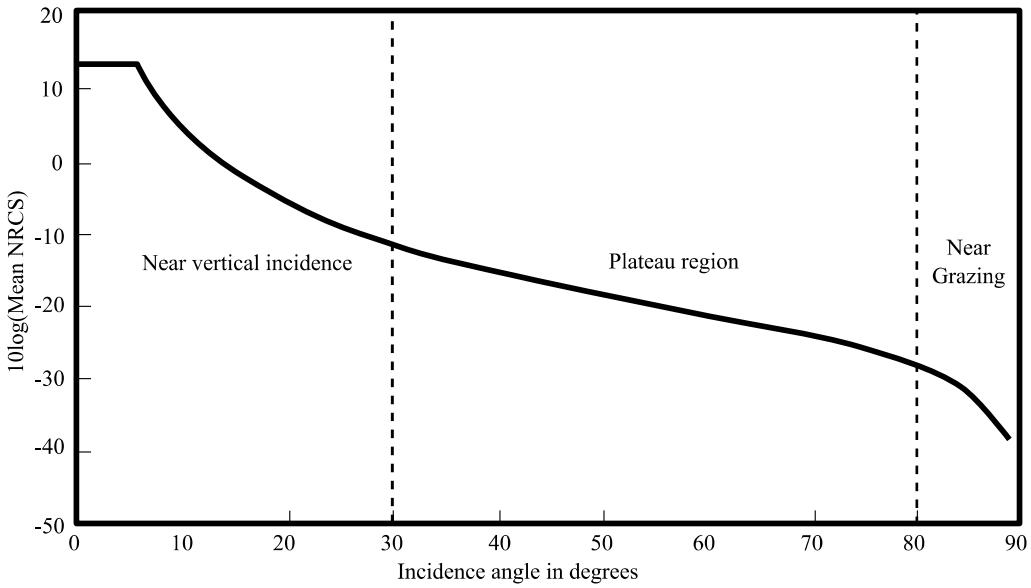

Figure 4: Mean normalised RCS for incidence angles between 0° and 90°. Figure adapted from Levanon [3], Figure 4.9, page 79.

Clutter has been recorded according to various conventions, and the most encountered is whether data is recorded as a function of incidence angle α_i, depression angle α_d or grazing angle α_g. Figure 5 shows the different angles. The incidence angle is the angle between the normal (of the surface below the radar) and the look direction of the radar. The depression angle is the angle between a horizontal line (parallel to the tangent of the surface) and the look direction. The grazing angle is the angle between the tangent of the earth's surface and the look direction. Grazing angle and depression angle have different numerical values when a spherical earth approximation has been made. For the flat earth approximation the grazing and depression angle are equal. The depression angle is also equal to 90° minus the incidence angle.

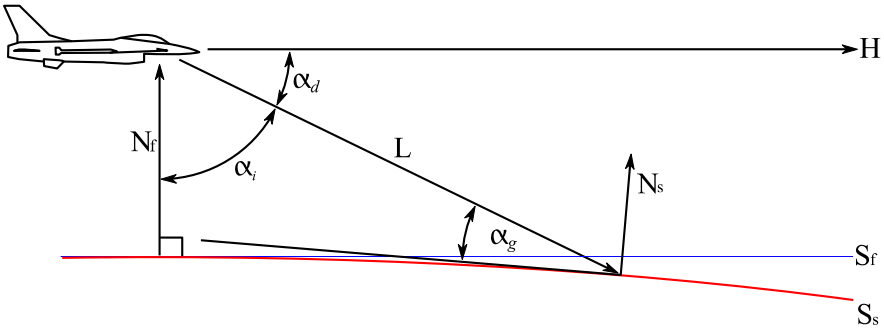

Figure 5: Definition of the independence angle, depression angle and grazing angle.

Depending on the range resolution of the radar and the incidence angle there are two different approximations for the surface geometry. The different surface geometries yield different results when the area is calculated. The first is the calculation when the range resolution is larger than the length of the total illuminated area of the main beam (4). This case usually occurs when radar is pointing relatively directly at the ground. When the radar is pointing in such a manner that the main beam grazes the surface, the range resolution has to be used as a limiting factor for the area calculation instead of the elevation beam width of the antenna. Equation (5) can be used to calculate this area [3].

$$S = \frac{R^2 \theta_H \theta_V}{\sin \alpha_d} \tag{4}$$

$$S = \frac{R \theta_H c t_p}{2 \cos \alpha} \tag{5}$$

Figure 6 shows the illuminated area of the ground with respect to the main beam width and the range resolution. The top panel is the side view of the geometry and the bottom panel is the top view of the geometry. The dark grey area represents the range resolution. The geometry is described by the height of the radar h, the vertical and horizontal beam widths θ_V and θ_H, the depression angle α_d, the range R, the range resolution t_p, and the speed of light c.

The separation of the surface area from the reflective properties of that surface allows the reflective properties of different surfaces to be characterised. This characterisation is usually grouped by different types of surfaces. The constant gamma (γ) model describes radar cross section with respect to the numerical value γ, and the grazing angle α_g (6). A different value for γ is chosen depending on the type of terrain. A few types of terrain and the corresponding value ranges for γ are listed in Table 1 for an X-Band radar.

$$\sigma^0 = \gamma \sin(\alpha_g) \tag{6}$$

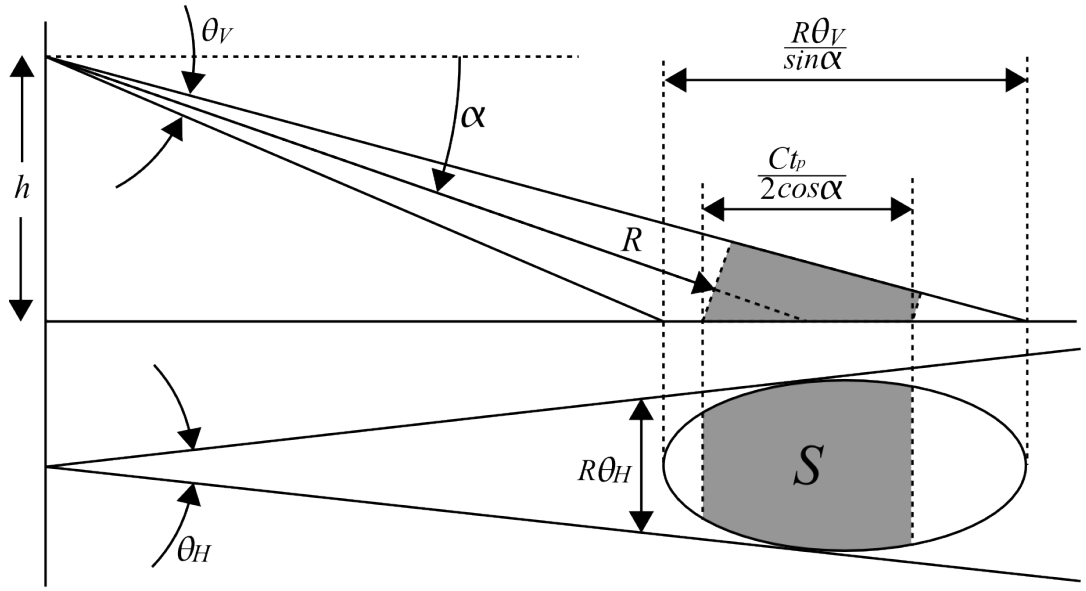

Figure 6: The side view and top view of the geometry for the calculation of ground clutter surface area. Adapted from Levanon [3], Figure 4.15, page 84.

Table 1: Typical values for the gamma parameter in the constant gamma model for different types of terrain at X-Band [2].

γ range (dB)	Type of terrain
-5	Urban or mountainous regions.
-10 to -15	Crops, bushes and trees.
-20	Desert, grassland and marshy terrain.

The constant gamma model only describes the mean value of the radar clutter. The model is a valid approximation for grazing angles between the critical angle α_c and 60° [4]. The critical angle is the lower limit where the constant gamma model can be considered valid and is dependent on the radar transmit wavelength λ and the root mean square (RMS) surface height deviation from the average σ_h as given in (7).

$$\alpha_c = \frac{\lambda}{4\pi\sigma_h} \tag{7}$$

Normal rolling terrain has a value of $\sigma_h = 15\text{m}$. The value for the critical angle is usually less than 5° [2]. The critical angle is less when the terrain roughness increases, or when a higher frequency is utilised by the radar.

A simple model like the constant gamma model is only an approximation to the actual behaviour of land clutter. Data from Long [5] shows measurements of terrain reflectivity for incidence angles between 10° and 80°. The measurements were presented in terms of γ model form. Figures 7 and 8 show that γ is anything but constant when measured over ploughed ground, disked ground, half

inch wheat, one and a half inch wheat, 15 inch grass, an asphalt runway and a concrete runway. Where the constant γ model identified crops, bushes and trees to be in the range of -10 to -15 dB, the data from Long shows that the range can be between -10 to -30 dB [5]. This wide range of terrain surface types with large ranges of values for gamma shows that the constant gamma model rarely holds true.

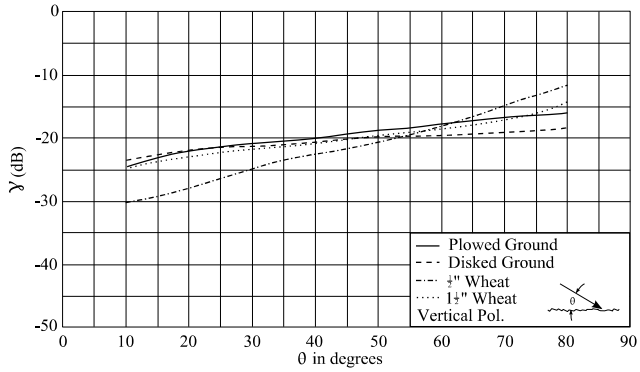

Figure 7: Comparison of reflected power vs. grazing angle for different types of terrain. Adapted from Long [5], Figure 2-11, page 232.

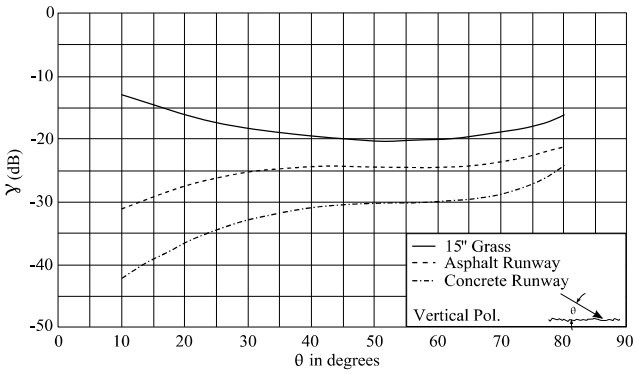

Figure 8: Comparison of reflected power vs. grazing angle for different terrain. Adapted from Long [5], Figure 6-10, page 231.

An added complication to clutter modelling is that the radar's echo from terrain is not necessarily present along the entire range line. The scatterers can be spatially distributed in a single range line as well as spatially distributed in the area surrounding the radar.

2.2 Spatial Distribution

The spatial distribution of ground clutter describes the statistical properties of the geographic positions of the clutter in the region surrounding the radar. This area is sensed by the radar in terms of range, azimuth and elevation. The spatial distribution of clutter is dependent on the position and height of the radar, as well as the terrain roughness [2]. Initially it was thought that clutter could be described as homogeneous with a sandpaper like roughness [6]. This was true for low resolution radars where clutter was sensed by the radar, all the way up to the radar horizon. As analogue to digital converter technology improved, it became possible to increase the radar's range resolution. With the increase in resolution separate clutter patches could be distinguished from each other. As a result the homogeneous assumption has since been found to be invalid for the higher resolution radars [6].

The increase in range resolution lead to clutter patches, and within these patches clutter was found to be spatially non-uniformly distributed. As the range resolution increases there is less averaging of clutter within the cells, and different independent features of the terrain become the dominant contributors within the clutter patch. The radar cross section of these independent features differ

widely from range cell to range cell. Figure 9 shows these effects for pulse lengths of 3.2 us, 0.8 us and 0.2 us [4]. This decrease in averaging affects the temporal variations of the clutter. The temporal distributions have been found to be Rayleigh distributed for larger (low-resolution) range cells gradually changing to Rician distributed statistics for smaller range cells. The change in distribution is due to the number of scatterers in a range cell. Many scatterers produce Rayleigh distributed backscatter. As the number of scatterers decrease, a few become dominant resulting in a Rice distribution [6].

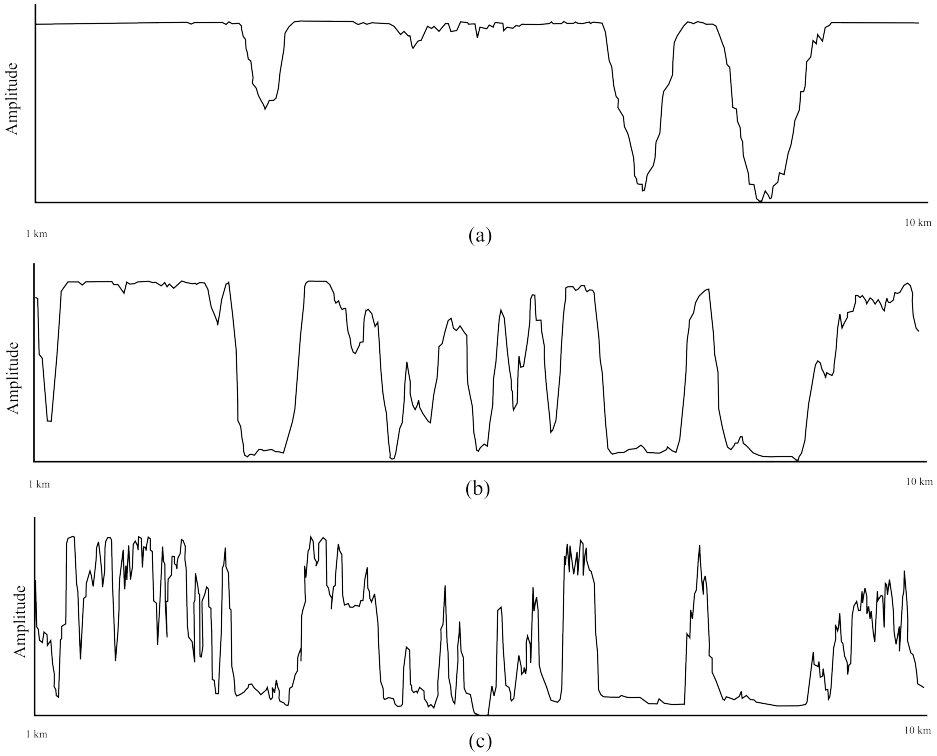

Pulse length set to (a) 3.2 us, (b) 0.8 us, and (c) 0.2 us.

Figure 9: Land clutter returns for varying pulse lengths. Adapted from Nathanson [4], Figure 7.26, page 330.

2.3 Amplitude Distribution

The amplitude distribution describes the variation in the magnitude of the clutter returns. A common mistake made when the amplitude distribution for clutter is discussed is to assume that this distribution is the distribution for a range cell, or a scatterer in the range line. This is not the case. The amplitude distributions for clutter are the distribution averaged over a clutter patch which can be the combination of many range cells. The temporal characteristics of clutter instead describe the amplitude distributions for single range cells [6]. The amplitude distribution measurement is a combination of the spatial and temporal characteristics of clutter.

To obtain the amplitude description of ground clutter, probability density functions are matched to the distributions of recorded data such as those found in Ulaby [7]. Figure 10 shows a histogram comparing the frequency of occurrence of the backscatter coefficient for Grasses measured at X-band with HH polarization at an incidence angle of 20°. A total of 377 samples were recorded from pulse to pulse (slow time). The 5% occurrence level, mean and 95% occurrence level are indicated by the arrows. This figure shows how the mean value of clutter power as described by the constant γ model, from Section 2.1, only accounts for a small part of the behaviour of the clutter. Figure 11 shows this effect across incidence angles between 0° and 80°.

To use the clutter results from literature, the radar system design, the signal processing and the data post processing has to be taken into consideration. Different literature sources may describe clutter with either Gaussian, or Rayleigh distributions, which are two completely different distributions. But when the background information is considered these two distributions could describe exactly the same behaviour for clutter. The Rayleigh distribution is the absolute value of the Gaussian distribution stemming from taking the magnitude of the complex input signals. The next sections describe various distributions which are encountered in clutter measurement literature.

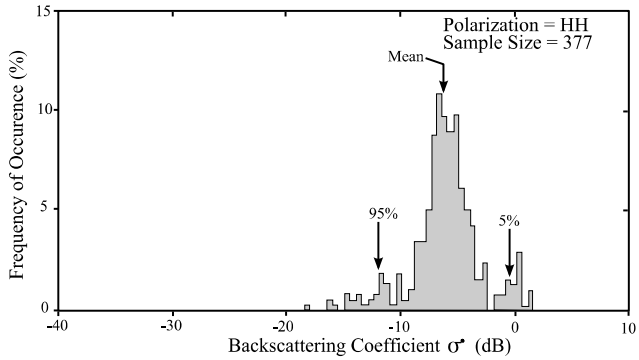

Figure 10: Histogram comparing the frequency of occurrence of the backscatter coefficient for Grasses. Adapted from Ulaby [7], page 185.

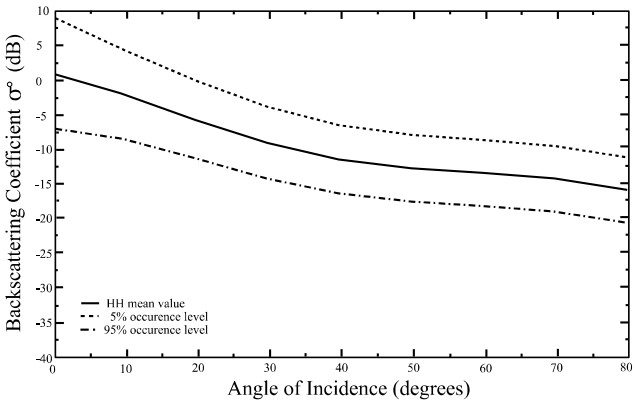

Figure 11: Plot showing the mean backscatter value for grasses. Adapted from Ulaby [7], page 184.

2.3.1 Gaussian (Normal) Distribution

The Gaussian (also known as the Normal, or bell curve) distribution, as used for clutter modelling, flows from the central limit theorem. The central limit theorem states that if independent and identically distributed (IID) random variables are summed their probability density function tends to a normal distribution as the number of random variables in the summation tends to infinity. Initially radars had poor range resolution and large stretches of land were thus contained in a single range cell. This poor resolution had a central limit theorem effect on the distribution of the clutter signal

as the sum of many independent scatterers was sensed in a single range cell, thus resulting in a normal distribution on I and Q. The PDF of the normal distribution is given in (8) with μ the mean, and σ the standard deviation [8].

$$P(x;\mu,\sigma)=\frac{1}{\sigma\sqrt{2\pi}}e^{\frac{-(x-\mu)^2}{2\sigma^2}} \tag{8}$$

2.3.2 Contaminated Normal Distribution

The contaminated normal distribution has been used for clutter modelling because some observations had heavy tails not described by the normal distribution. A contaminated normal distribution is a distribution where samples are drawn from two different normal distributions, to obtain heavy tails. The majority of the samples are drawn from one normal distribution with mean μ_1 and standard deviation σ_1. The remaining samples are drawn from a second normal distribution with a mean μ_2 and a standard deviation σ_2. The relationship between the two distributions are an equal mean $\mu_1=\mu_2$, and a larger variance for the less sampled distribution $\sigma_1 \ll \sigma_2$. This produces a heavy tailed, contaminated normal distribution [9].

2.3.3 Rayleigh Distribution

The Rayleigh distribution is used in clutter modelling for scenarios where there are large incidence angles, such as for the airborne radar case. Traditionally the Rayleigh distribution is used to model noise, and stems from envelope detection of radar pulses in a Gaussian noise environment [10]. The Rayleigh random variable (RV) is related to the Gaussian RV by (9) where g_1 and g_2 are independent Gaussian RVs with zero mean and equal variance.

$$r=\sqrt{g_1^2+g_2^2} \tag{9}$$

The PDF of the Rayleigh distribution is given by (10) with x the statistical instantaneous amplitude, and α the scale parameter of the distribution. The parameter α modifies the scale of the Rayleigh distribution. A smaller α results in a shorter and more peaked distribution, while a larger α results in a longer more flat distribution. Deriving (10) from (8) and (9) one observes that the standard deviation of the Gaussian random variables, become the scale parameter α of the Rayleigh distribution. The mean μ and variance are given by (11) and (12) respectively [11].

$$P(x;\alpha)=\frac{x}{\alpha^2}e^{\frac{-x^2}{2\alpha^2}},\ x\geq 0 \tag{10}$$

$$\mu=\alpha\sqrt{\frac{\pi}{2}} \tag{11}$$

$$var(x)=\left(\frac{4-\pi}{2}\right)\alpha^2 \tag{12}$$

The higher order moments of a distribution are often used to match recorded data to a specific PDF. The order of the moment is denoted by the variable n in (13).

$$\langle x^n\rangle=2^{n/2}\alpha^n\Gamma\left(1+\frac{n}{2}\right) \tag{13}$$

2.3.4 Log-Normal Distribution

The Log-Normal distribution is a distribution where the log of the distribution is normally distributed. This distribution is sometimes found for clutter at low incidence angles, but is more generally used to describe the spatial distribution of ground clutter. The log in this case is the natural logarithm. The PDF of the Log-Normal distribution is given in (14), with μ the mean and σ the standard deviation. The moments of the Log-Normal distribution are given by equation (15) [11], [12].

$$P(x;\sigma,\mu)=\frac{1}{x\sigma\sqrt{2\pi}}e^{-\frac{(\ln(x)-\mu)^2}{2\sigma^2}},\ x\geq 0 \tag{14}$$

$$\langle x^n\rangle=e^{n\mu+\frac{n^2\sigma^2}{2}} \tag{15}$$

2.3.5 Exponential Distribution

Extensive tests focusing on the accuracy of clutter measurements at the Massachusetts Institute of Technology (MIT) Lincoln laboratory during the 1990s lead to the exponential model. The exponential distribution PDF is given in (16). The variable r is called the rate parameter of the exponential distribution. A comparison between the Gaussian and exponential models have less than a factor two difference in the tails, which is negligible for scanning radars [8], [13]. Windblown clutter has been found to be a better match to the exponential distribution than the Gaussian distribution [14], [15]. The exponential model as described here is on an amplitude scale.

$$P(x;r)=re^{-rx} \tag{16}$$

2.3.6 K-Distribution

The K-Distribution has been found to be a relatively accurate model for sea clutter, and is sometimes also appropriate for describing ground clutter. The K-Distribution has a scale and shape parameter, b and v respectively. $\Gamma(v)$ is the gamma function and $K_v(x)$ is the modified Bessel function. Similarly to the Rayleigh distribution the scale parameter condenses or elongates the PDF,

but does not alter the shape. The shape parameter alters the shape of the distribution. For a large value of the shape parameter, $v \to \infty$, the PDF becomes Rayleigh. The PDF and moments of the K-Distribution are given in equations (17) and (18) [12], [16].

$$P(x;b,v)=\frac{4b^{(v+1)/2}}{\Gamma(v)}x^{v}K_{v-1}(2\sqrt{b}x),\ x\geq 0 \tag{17}$$

$$\langle x^{n}\rangle=b^{-n/2}\frac{\Gamma(1+n/2)\Gamma(v+n/2)}{\Gamma(v)} \tag{18}$$

2.3.7 Weibull Distribution

For certain conditions the Weibull distribution has been found to be the most ideal fit to radar clutter. The Weibull distribution is capable of changing the size of its tail, by changing the shape parameter k. The variable λ is the scale parameter of the distribution. The PDF of the Weibull distribution is given in (19) and the moments in (20). The Weibull distribution is easier to calculate than the K-Distribution, and is generally a more accurate fit for clutter than the Log-Normal distribution. Setting the shape parameter $k=2$ the PDF reduces to that of a Rayleigh distribution. This versatility makes the Weibull distribution a popular choice for clutter modelling [17].

$$P(x;\lambda,k)=\frac{k}{\lambda}\left(\frac{x}{\lambda}\right)^{k-1}e^{-\left(\frac{x}{\lambda}\right)^{k}},\ x\geq 0 \tag{19}$$

$$\langle x^{n}\rangle=\lambda^{n}\Gamma\left(\frac{n}{k}+1\right) \tag{20}$$

2.3.8 Alpha Stable Distribution

The alpha stable distribution is the latest distribution to be used in the continuing process of describing clutter statistics with more general distributions. The alpha stable distribution densities have two important properties that make them ideal for clutter modelling. The first property is the stability property, which states that the weighted sum of independent alpha stable random variables is again stable, with the same characteristic exponent α. The second property is the generalised central limit theorem. Unlike the Normal distribution an alpha stable distribution can also be asymmetric about the mean. The alpha stable distribution is given in (21). The characteristic exponent α is bounded between 0 and 2. The variable δ is the location parameter and γ is the dispersion parameter [18]. This distribution has not yet found widespread use in the literature on radar clutter.

$$P(x;\delta,\gamma,\alpha)=e^{j\delta x-\gamma|x^{\alpha}|} \tag{21}$$

Each of the distributions mentioned in sections 2.3.1 to 2.3.8 have a place in describing clutter. The three different PDFs that are the most frequently used when describing random nature of ground clutter returns are the Rayleigh distribution, Log-normal distribution, and Weibull distribution. Figures 12 to 15 show these distributions. For the Log-Normal distribution the offset parameter μ was held constant at 0, and for the Weibull distribution the scale parameter λ was held constant at a value of 1.

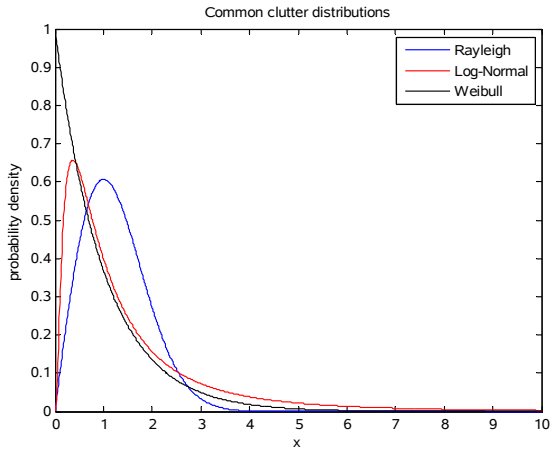

Figure 12: A comparison of the Rayleigh, Log-Normal and Weibull distributions.

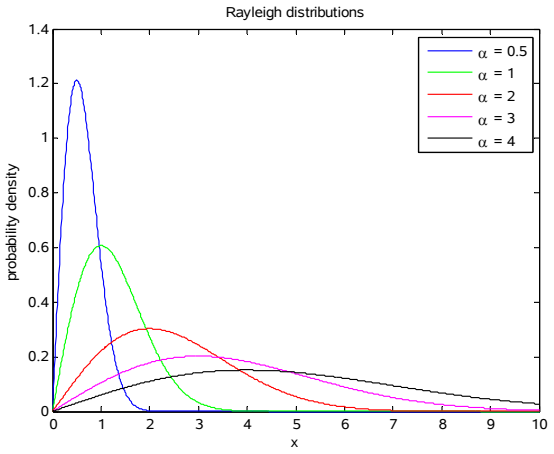

Figure 13: The Rayleigh distributions with different values for the scale parameter alpha.

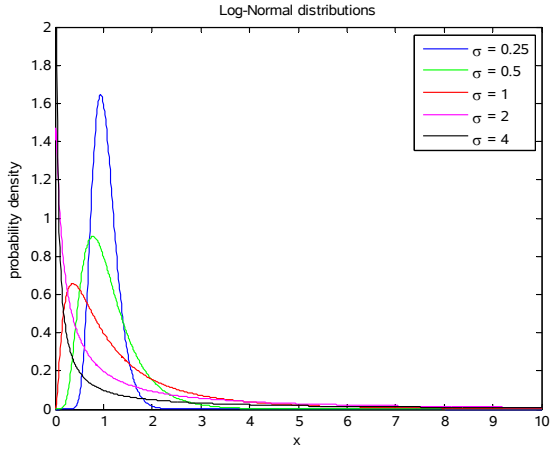

Figure 14: The Log-Normal distributions with different values for the scale parameter sigma.

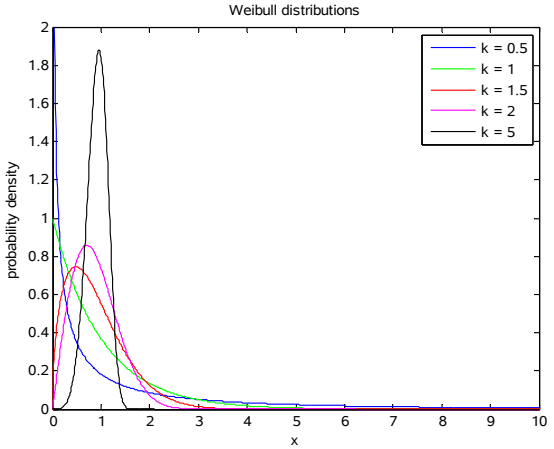

Figure 15: The Weibull distributions with different values for the shape parameter k.

The clutter PDF is usually assumed to be Rayleigh when simplifying assumptions are made. The Weibull and Log-Normal distribution are accepted PDFs to use for ground clutter. The K-Distribution has been fitted to land clutter, but is more commonly used to describe sea clutter. Both the Weibull distribution and the K-distribution reduce to a Rayleigh distribution with the correct input parameters. This makes them better for radar clutter modelling since they incorporate both the low and high resolution clutter scenarios. Although Log-Normal distributions have been fitted to

radar data, the inability of this distribution to reduce to a Rayleigh distribution makes them less attractive for radar clutter simulation [19]. The simplest "realistic" simulation possible is that of Rayleigh distributed clutter back-scatter.

The above mentioned distributions result after the detection algorithms are applied to the input signal as part of the radar signal processing. The amplitude detection process of the radar transforms symmetrical PDFs found on the in-phase and quadrature channels into these asymmetrical PDFs. The amplitude is obtained by calculating the magnitude of the in-phase and quadrature channels of the received signal of the radar as shown in equation (22) [20]. The next section discusses the spectral properties of radar clutter.

$$magnitude = \sqrt{I^2 + Q^2} \tag{22}$$

2.4 Spectral Distribution

For this discussion of the spectral characteristics of clutter a pulsed waveform system is assumed. A rectangular pulse train will have spectral lines with a $sinc(x)$ envelope. The $sinc(x)$ function is given in equation (23). The shape of the $sinc(x)$ envelope is determined by the transmitted pulse width t_p, with the first nulls occurring at $f_o \pm 1/t_p$ as shown in Figure 16. The symbol f_o is the centre frequency of the radar. The spacing of the spectral lines is equal to the pulse repetition frequency (PRF) of the radar. These spectral lines are widened by the motion of clutter, and by the movement of the radar due to scanning or radar platform motion.

$$sinc(x) = \frac{\sin(\pi x)}{\pi x} \tag{23}$$

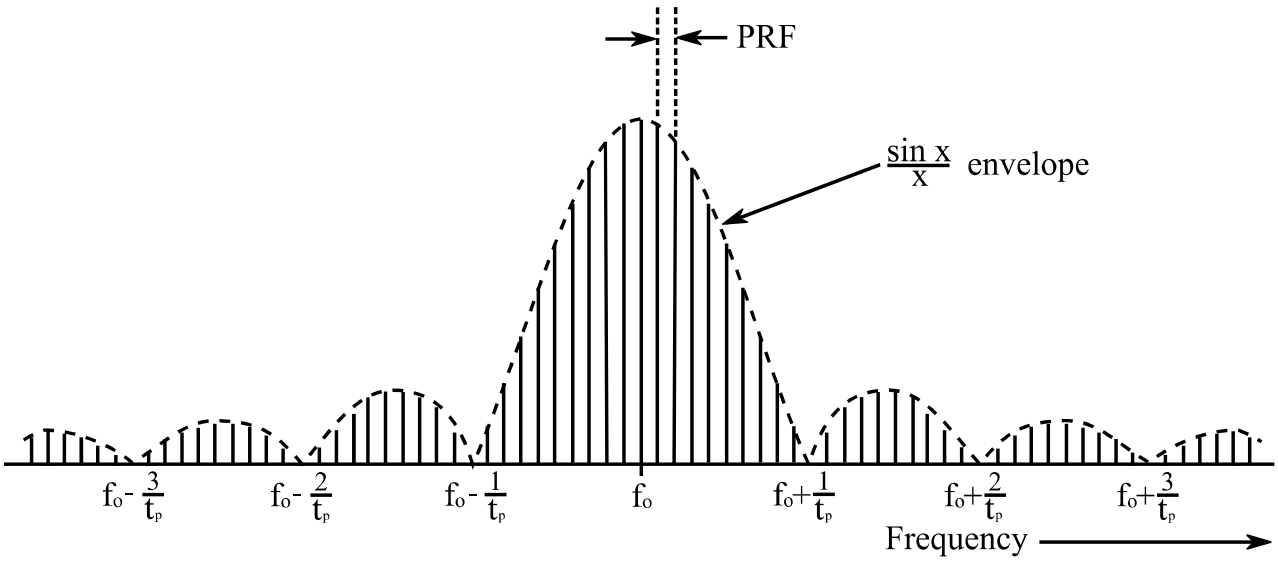

Figure 16: Frequency spectrum of a pulsed signal showing the effects of the pulse length and the PRF of the radar. Adapted from Skolnik [21], Figure 2.11, page 2.11.

Equation (24) is used to calculate the Doppler frequency shift f_d based on the wavelength λ of the radar and the angular velocity of the object v at an angle θ [21].

$$f_d = \frac{2v\cos(\theta)}{\lambda} = \frac{2v\cos(\theta)f_o}{c} \tag{24}$$

The Doppler spectrum of a single stationary point target is a sharp spike at zero Doppler shift, as seen by the measurement of a water tower at 6.4 km using the Phase One radar (Figure 17). A small amount of ground clutter from trees is visible at the base of the spike close to the noise level of the radar. The narrow nature of the tree clutter is due to a lack of wind and thus the trees are relatively stationary. The stationary nature of the trees means there is a lack of non-zero Doppler components. When the wind is present the non-zero Doppler components of the clutter become more prominent as seen for different wind speeds through trees measured with the LCE radar at 7.0 km (Figure 18). As the wind speed increases the width of the spectral shape of the clutter becomes wider. The spectral shape of the clutter is essentially independent of frequency and polarization. The frequency independence of clutter is unexpected because main scatterers differ between VHF (large branches) and X-band (leaves and twigs). The frequency and wavelength response is further investigated in Section 2.6.

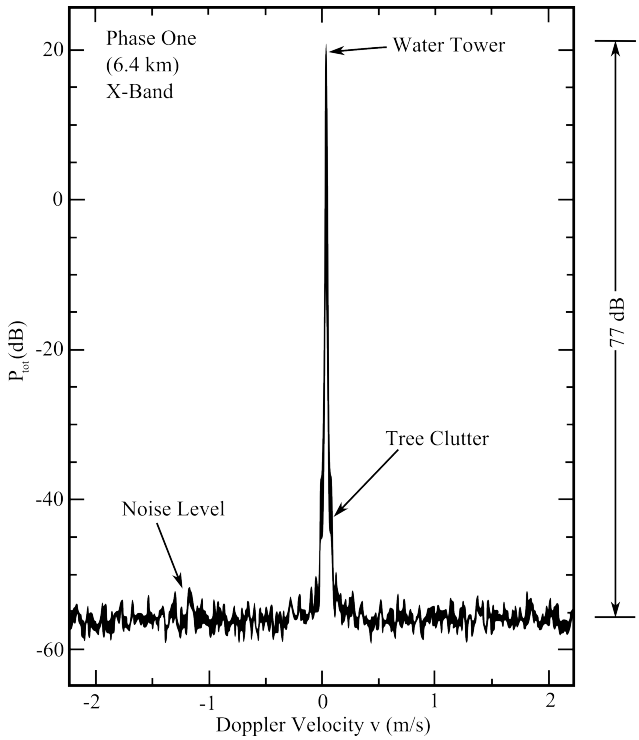

Figure 17: Spectral shape for a single target. Adapted from Billingsley [6], Figure 6.4, page 587.

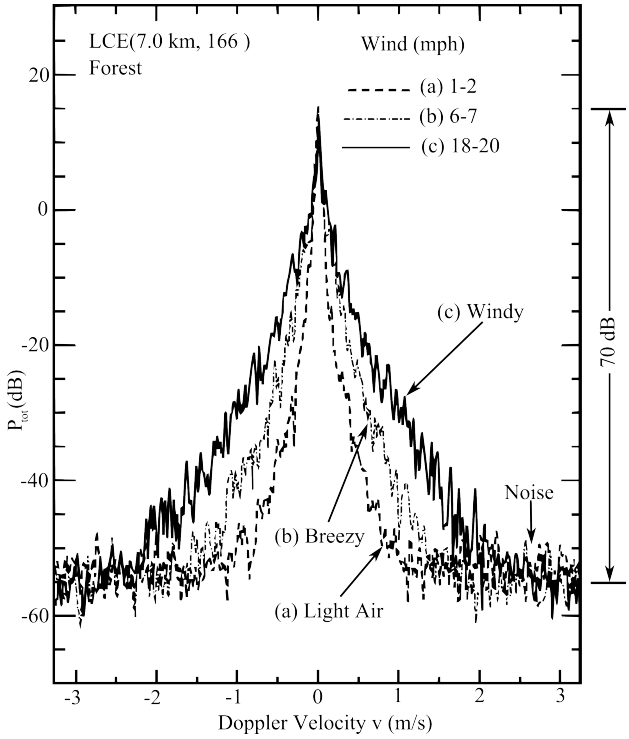

Figure 18: Spectral shape for windblown trees. Adapted from Billingsley [6], Figure 6.8, page 593.

Billingsley [6] used an exponential decay model for the spectrum shape of ground clutter. The total power is described by equation (25). This analytical expression describes both the alternating current (AC) power of the clutter and the direct current (DC) power of the clutter. Doppler modulation is responsible for the AC power $P_{ac}(v)$, and stationary objects are responsible for the DC power described with the delta Dirac function $\delta(v)$. An AC to DC ratio is used to describe the distribution of the power between those two parts of the equation. The AC power is modelled as an exponential decay (26). Table 2 lists common values for β for different wind speeds. To compare the model with measured results the values generated by the model have to be reduced by $10\log_{10}(r+1)$, where r is the ratio of DC to AC power. The value for beta varies almost linearly with wind speed as described by equation (27), where w is the wind speed. This equation slightly underestimates β for light wind speeds and slightly overestimates β for gale force winds.

$$P_{tot}=\frac{r}{r+1}\delta(v)+\frac{1}{r+1}P_{ac}(v),-\infty<v<\infty \tag{25}$$

$$P_{ac}(v)=\frac{\beta}{2}\exp(-\beta|v|),-\infty<v<\infty \tag{26}$$

$$\beta^{-1}=0.1048[\log_{10}w+0.4147] \tag{27}$$

The AC to DC ratio is dependent on the radar's RF frequency and can be determined using equation (28) for radar range cells which contain windblown trees where w is the wind speed in miles per hour and f_o is the radar carrier frequency in GHz [6]. The AC to DC ratio can also give an indication of the correlation of the signal since a smaller AC component indicates a smaller spectral bandwidth and thus more correlation from pules to pulse. The next section discusses the correlation properties of radar clutter.

$$r=489.8\,w^{-1.55}f_o^{-1.21} \tag{28}$$

Table 2: Exponential AC Shape Parameter for certain wind speeds [6].

Wind Conditions	Wind Speed (mph)	Exponential AC Shape Parameter β $(m/s)^{-1}$	
		Typical	Worst Case
Light air	1-7	12	-
Breezy	7-15	8	-
Windy	15-30	5.7	5.2
Gale force (estimated)	30-60	4.3	3.8

2.5 Correlation Properties

The correlated nature of ground clutter can greatly increase a radar's false alarm rates. It is therefore important to correctly model the correlation properties of ground clutter. If the correlation is modelled correctly, the radar's clutter cancellation algorithms can be tested correctly. Radar clutter can be correlated spatially from range cell to range cell (fast time), or temporally from pulse to pulse (slow time). Correlation in azimuth is also measured from pulse to pulse.

Spatial correlation is defined in both the range and azimuth directions. The correlation in space for clutter is dependent on the specific scatterers in each range cell, and since the range cell contains unique scatterers (in other words a single scatterer is not present in more than one range cell) clutter decorrelates within the extent of a single range cell. However it has been found that at times slightly higher correlation is found at higher angles (angles not defined as low incidence angels) over uniform landscapes. Similarly if the beam of the radar moves enough in azimuth that it no-longer illuminates the same azimuth cell, the contents of that cell is unique so that clutter decorrelates within a single range cell. The decorrelation within a single range cell, allow for the assumption that clutter is spatially uncorrelated, the extent of which is shown in Figures 19 and 20 [6].

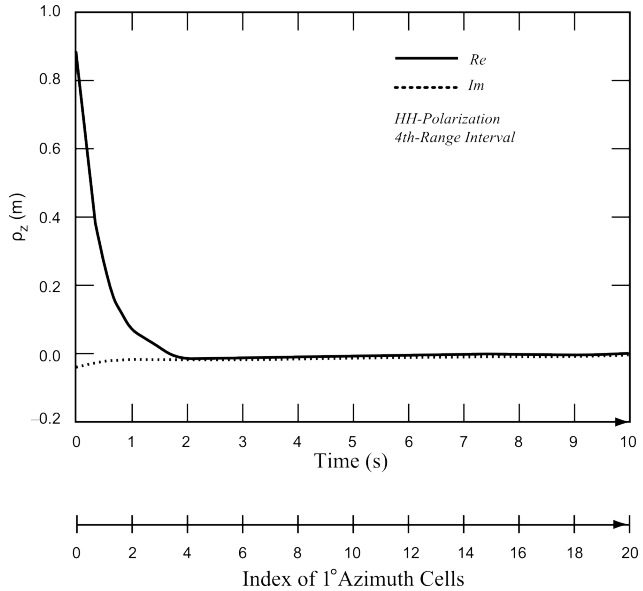

Figure 19: Correlation coefficient in the azimuth dimension measured in slow time. Adapted from Billingsley [6], Figure 4.26, page 341.

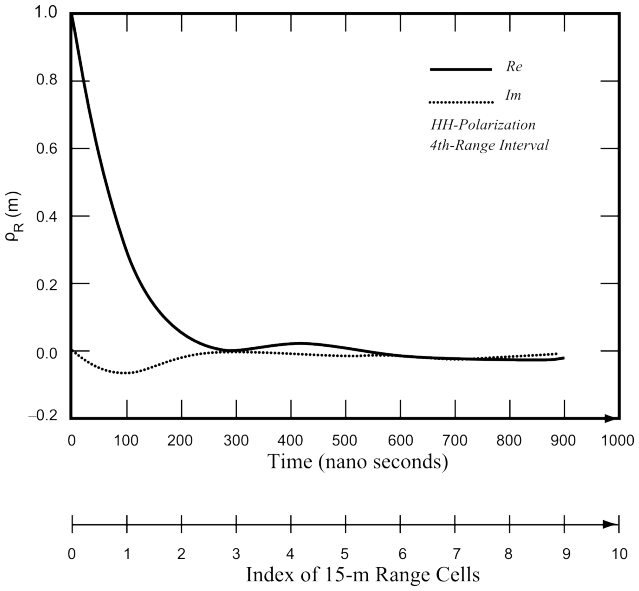

Figure 20: Correlation coefficient in the range dimension measured in fast time. Adapted from Billingsley [6], Figure 4.27, page 342.

Temporal correlation of clutter is an alternative description of the spectral content of the same clutter signal. The normalised autocorrelation function can be used to calculate the correlation time for a signal. The correlation time indicates the extent to which the clutter is correlated. The correlation times of clutter signals are inversely proportional to the centre frequency of the transmitted signal when only considering the Doppler equation and keeping all other variables

constant. The actual trend deviates slightly from this because correlation times are dependent on the scattering characteristics of the environment, and different frequencies interact in different ways with the environment. The correlation times are also inversely proportional to the wind speed which translates to a longer decorrelation time for slower wind speeds. Table 3 shows correlation times as the signal's correlation reaches half $\tau_{1/2}$, or one over e, $\tau_{1/e}$ of its peak value, for different frequency bands [6]. The next section discusses the wavelength and polarization properties of clutter.

Table 3: Correlation times for radar returns from windblown trees on a windy day [6].

Frequency Band	Correlation Time (s)	
	$\tau_{1/2}$	$\tau_{1/e}$
VHF	4.01*	5.04*
UHF	0.69	0.94
L-Band	0.67	0.95
S-Band	0.062	0.081
X-Band	0.033	0.490

Values with a * have been estimated by extrapolation.

2.6 Wavelength and Polarization properties

According to [4] the effect of the radar transmit frequency, on the backscatter coefficient as a function of grazing angle, is a near linear increase on average. Wavelength is inversely related to frequency and thus an inverse relationship between the radar wavelength and clutter backscatter is observed. The complex interactions between environmental factors and the electromagnetic waves complicates the search for a clear relationship. Different environments interact differently with the wavelength of the radar pulses [4], [5]. Figure 21 show the general trends in land clutter backscatter for rural land for different wavelengths.

Differences in clutter backscatter for different polarizations are due to the physical structure of the terrain. Tall grass is aligned with vertical polarization and as a result will reflect more energy than a horizontally polarized wave. If the grass is dry, this effect will be less. For large areas of terrain horizontal polarization tends to be a few decibels less than vertical polarization between grazing angles of 1° and 3°. The largest factor influencing polarization sensitivity is the alignment of scattering objects with the polarization direction, and multipath effects. Large flat surfaces such as bodies of water or highways also contribute to the polarization sensitivity. Polarization generally

has little effect on the backscatter coefficient of ground clutter, and even though individual measurements can show a difference in clutter backscatter, there is no clear trend from the measured data in literature [4], [5].

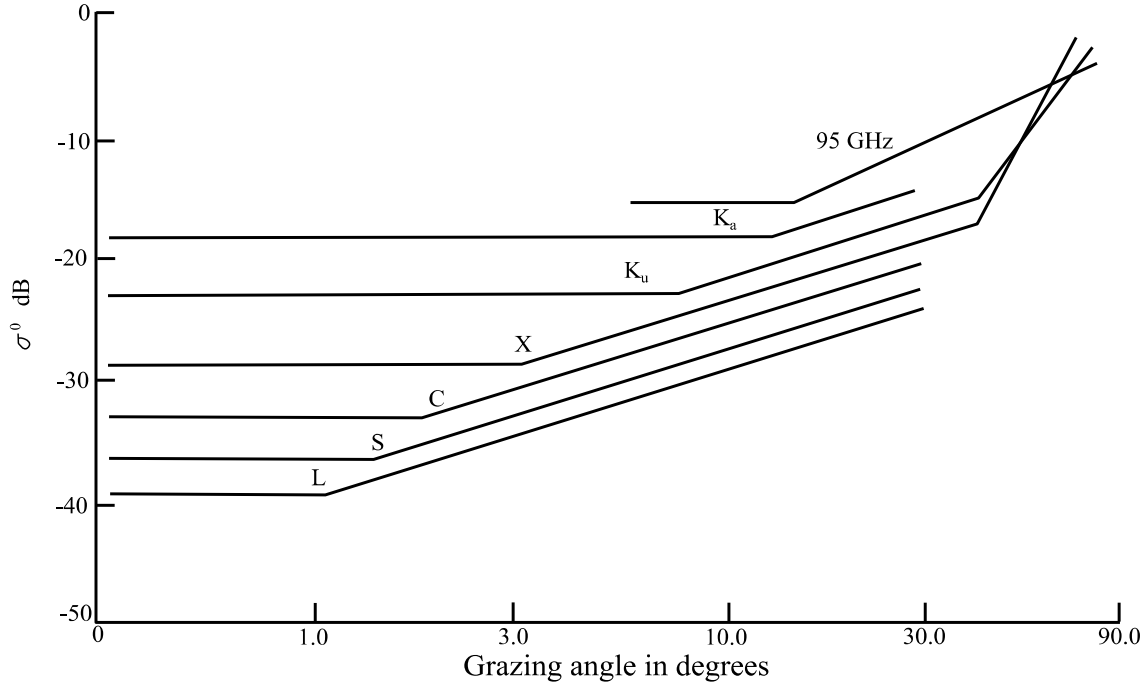

Figure 21: Clutter average power for different grazing angles and radar frequency bands for a rural environment. Adapted from Nathanson [4], Figure 7.23, page 323.

2.7 Conclusion

There are many parameters that describe the characteristics of radar clutter which have to be taken into account when radar environment simulation is considered. For ground clutter the clutter statistics are the most important aspect to focus on for radar environment simulation.

For the rest of the study, the scenario with the lowest number of variables possible will be assumed to simplify certain aspects of the clutter simulation. An increasing radar cross section results in a simple magnitude increase in the received clutter signal. The wavelength and polarization can similarly be adjusted and compensated for when the radar environment is stationary within a constant look direction. The spatial distribution is a matter of enabling clutter for certain parts of the range line and disabling them for others and is thus not very difficult to simulate to the required detail. The amplitude distribution, and Doppler spectrum of the clutter are the main aspects that have to be modelled for clutter simulation.

The next section of the report will investigate methods for simulating clutter in software. The software simulation will determine the extent to which realistic clutter simulation is possible.

3 SIMULATION OF CLUTTER, IMPLEMENTED FROM LITERATURE

To simulate realistic clutter for a radar the entire sequence of events that produces the final result has to be taken into account. The typical sequence of events for a mono-static radar system (as shown in Figure 22) starts with the transmission of an electro-magnetic pulse. The transmitter and receiver for a mono-static radar is co-located, which means the propagation path to and from a clutter patch is the same. This is not the case for bi-static radar, where the paths may have completely different propagation effects. The transmitted radar pulse propagates through the air which is the transmission medium of the electromagnetic energy. The radar pulse then interacts with the environment. Some of the energy of the pulse is absorbed and some is reflected. The reflected energy is scattered in many directions, some back in the direction of the radar. The radar receives the reflected energy, and senses the environment by performing radar signal processing on the measured reflected electro-magnetic energy. The measured reflected energy is responsible for the ability of the radar to identify the object(s) of interest. The measured energy reflected from the environment is identified as clutter, and is usually not of any value.

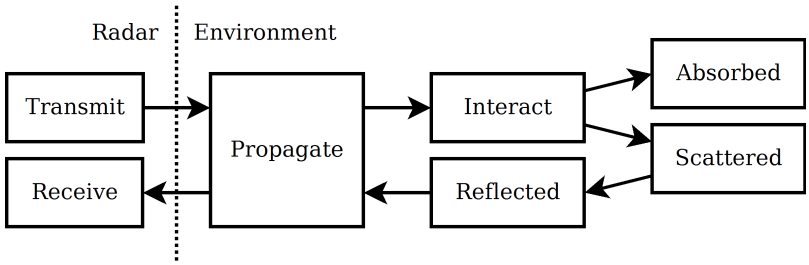

Figure 22: Functional block diagram of the simulation of mono-static radar clutter.

Radar signal processing in its simplest form is a multi step process consisting of down-mixing, Hilbert-filtering, and amplitude detection, also known as quadrature down-mixing. Doppler processing uses a discrete Fourier transform (DFT) to also determine the Doppler shift caused by the motion of objects in the environment. The challenge for the radar clutter simulator is to produce realistic clutter after the radar signal processing has been performed, by injecting a signal between the transmission and reception of the radar pulse.

In the hypothetical case where hardware systems have infinite resources, any clutter statistics can be created if an appropriate algorithm can be identified. This ideal case is however not the reality and generating clutter is resource intensive. The goal of this study is to find clutter generation

algorithms that are as realistic as possible but at the same time implementable on a hardware platform with limited resources in real time. Because clutter is statistical in nature, at the core of the clutter generation algorithms are random number generators.

3.1 Random number generator

One way to generate a random number with a given PDF is to start with a random variable with a uniform PDF. A transformation is then applied to obtain a random variable with the desired distribution. There are four types of transformations that can be applied to a uniform random variable to obtain Gaussian random variables. The four methods are the CDF inversion method, transformation method, rejection method, and the recursive method [22].

The CDF inversion method uses the inverted CDF of any random variable, and is not limited to a Gaussian distribution. By definition the area underneath the PDF is equal to one. The range of the PDF can potentially span from $-\infty$ to ∞. Because of the unity area property, the CDF of a random variable can never be less than zero or more than one. The inverted CDF then can only have input values (x) in the range of 0 and 1. A generated uniform random number is in that same range, or can be transformed to be in that range by scaling and shifting operations, for use as the input values. Thus a uniform random number sequence can be used as the variable in the inverted CDF function and the resulting random variable sequence will take on the corresponding PDF [22], [23].

Transformation methods such as the Box-Muller transform [22], the central limit theorem, piecewise linear approximation using triangular distributions [24] and the Monty Python method [25]. These methods start with a uniform random number and perform a combination of mathematical operations to obtain another distribution. This is similar to the CDF inversion method, but the mathematical operations are not directly derived from any PDF, and they instead rely on other properties, for example, the central limit theorem. The central limit theorem is limited to generating Gaussian random variables, but can do so with low complexity and computational expense. The piecewise linear approximation with triangular distributions is not an exact method, but its strong point is that it not limited to a Gaussian distribution [23]. Both the Box-Muller and Monty Python are limited to a Gaussian distribution [22].

Rejection methods contain a conditional step in the generation of the random variable. This conditional step decides if generation of the random variable should continue, or if the current value at this point in the algorithm must be rejected. If the current variable is rejected a new variable is generated from scratch. This conditional step is not favourable for hardware implementation where a constant flow of new random variables is required [22].

The recursion method, also known as the Wallace method uses the property that a linear combination of Gaussian distributed random variables will produce a new Gaussian distributed random variable. The Wallace method uses matrix multiplications with Gaussian distributed random numbers to obtain a set of output Gaussian distributed random numbers. Because the previous numbers are used to generate new ones, there is correlation between old and new values. The Wallace method repeats the multiplication step a set number of times to sufficiently decorrelate the numbers before they are used [22]. The randomness of nature visible in clutter is not uncorrelated, and the degree of correlation is dependant on the environment. The Wallace methods lends itself to this requirement, but a method to reliably control the degree correlation is lacking, and the fact that numbers are generated in batches counts against its use for real time clutter simulation.

Radar clutter is described by both its amplitude statistics, and its autocorrelation function. The statistics are usually only valid for a single range cell. The autocorrelation is a measure of the correlated nature of the radar clutter. The autocorrelation can be obtained from a cross correlation of the radar backscattered signal with itself. The samples of the signal are obtained from pulse to pulse (slow time), from the same range cell. The spectrum is estimated by taking a DFT across multiple pulses. Different range cells in the radar pulse have to be investigated individually because the clutter statistics might vary between them due to spatial variability of clutter. Clutter is described as being spiky in nature because on average there is little variation in the return of a clutter patch, with large clutter returns occurring infrequently. The frequency of occurrence of the large clutter returns are described by the PDF of the clutter.

Clutter can be approximated using several different PDFs as described in Section 2.3, and methods exist that would allow for the simulation of nearly any distribution imaginable. The inverse cumulative density function method is one such an example, where if the cumulative density can be found (even if only numerically) the approximate distribution can be generated. Approximations are not a major cause for concern when clutter simulation is considered since the distributions used to model clutter have been fitted to recorded data and are thus themselves an approximation. Clutter simulation is of course different from standard random number generation in the sense that it is correlated. The correlated Gaussian random variable is at the core of generating correlated random variables having non-Gaussian distributions.

3.2 Correlated Gaussian Random Variable

Clutter is correlated from pulse to pulse (slow time). An uncorrelated random number will have a white spectrum which indicates it is uncorrelated from one value to the next, and thus not consistent with the behaviour of clutter. Linear system theory states that a signal with spectrum S_{in} can be filtered to produce a signal with a spectrum of S_{out} by passing it through a filter with a transfer function $H(f)$ as shown in equation (29). Thus filtering a white noise signal will correlate that signal.

$$S_{out}(f)=S_{in}(f)\left|(H(f))^2\right| \tag{29}$$

A correlated Gaussian random variable can be generated by introducing a dependency on the previous random sample, when generating the next one. Correlation can be introduced by filtering a Gaussian random variable. The lowest complexity digital filter would be a filter with a single feedback stage, for example a single pole infinite impulse response (IIR) filter which is a good approximation to the spectrum of measured ground clutter. The filtered spectrum is obtained using equation (30) [12]. This correlation is responsible for the shape of the spectrum of the clutter. The bandwidth of the clutter is the frequency at which the spectrum is three decibels (dB) less than the maximum value. The bandwidth of the filter is set to be equal to the required bandwidth of the clutter. A single pole infinite impulse response filter will produce a spectrum with an approximately Gaussian shape.

$$x_n=\eta x_{n-1}+\beta g_n \tag{30}$$

It is useful to obtain a relation between η, β and the -3 dB bandwidth of the filter. Taking the z-transform of equation (30), gives (31), which can be manipulated to give the transfer function (32).

$$Y(z)=\eta Y(z)z^{-1}+\beta X(z) \tag{31}$$

$$|H(z)|=\frac{Y(z)}{X(z)}=\frac{\beta}{1-z^{-1}\eta} \tag{32}$$

Adhering to the notation of the z-transform, the input variable g_n is represented by $X(z)$ and the output variable x_n is represented by $Y(z)$. For the filter to exhibit unity power gain the relation $\beta=\sqrt{1-\eta^2}$ has to hold [12]. The z-transform states that $z=Ae^{-j\phi}$, where A is the magnitude, j the imaginary unit, and ϕ the phase. Substituting $Ae^{-j\phi}$ for z, with A as unity and β with $\sqrt{1-\eta^2}$ into equation (32) produces equation 33.

$$|H(\phi)|^2=\frac{1-\eta^2}{1-2\eta\cos\phi+\eta^2} \tag{33}$$

The result of this is the ability to calculate the amplitude of the spectrum of the correlated Gaussian signal at any frequency and for any value of η. Because we are interested in controlling the -3 dB bandwidth, we have to calculate the position in the z-transformed space (which relates to a frequency) where the amplitude of the spectrum is half the value of the maximum amplitude, given in equation (34). Solving this equation leads to a relationship between η and ϕ given in equation (35).

$$\frac{|H(0)|^2}{2} = |H(\phi)|^2 \tag{34}$$

$$\eta = 2 - \cos\phi \pm \sqrt{(2-\cos\phi)^2 - 1} \tag{35}$$

The value for the phase ϕ relates to the required bandwidth of the filter with equation (36), where BW is the bandwidth and Fs is the sampling frequency both in Hertz (Hz).

$$\phi = 2\pi\left(\frac{BW}{Fs}\right) \tag{36}$$

This design equation gives the ability to design a filter with adjustable bandwidth, and correlated output with a Gaussian distribution. The performance of this filter design has to be evaluated with regards to frequency spectrum and probability density function.

3.3 Evaluation of clutter generators

To analyse the performance of a random variable generator, the error between its output and a theoretical baseline has to be calculated. At the most basic level, one has to compare the results from a random variable generator to those observed from actual measurements presented in the literature. In Section 2 a literature survey was presented and gave guidelines to the different aspects of clutter simulation that have to be verified and used for comparison. An aspect that adds an additional level of complexity to the evaluation of clutter generators is that the measured results do not fit the models that were drawn from the measurements perfectly. When simulating clutter the aim is to recreate the behaviour of clutter, but this is done by recreating the behaviour of a model that represents the clutter.

The evaluation methods chosen for this project were the analysis of the frequency spectrum and autocorrelation function. Signal theory states that the frequency spectrum and the autocorrelation function, are a Fourier transform pair [26]. Both these have been analysed because the literature on clutter does not always focus on either one, or both. Both the frequency spectrum and autocorrelation function describe the correlation behaviour of clutter. The other aspect of clutter is the amplitude distribution. There are various methods to investigate how close a generated

amplitude distribution comes to the theoretical one. For this project the χ^2 test and KL divergence value, have been used in conjunction with the calculation of the higher order moments to analyse how well a generated result matches that of the analytical model. To make the results more tractable and easier to interpret, distance methods to error calculation have also been employed.

The Kullback-Liebler (KL) divergence, often also referred to as the KL distance is a statistical measure of the distance between two distributions. A statistical distance measure like the KL distance is a measure of how alike two distributions are. A smaller distance value implies a smaller difference between two distributions, while a larger distance implies a larger difference between two distributions. The KL distance between distribution p and distribution q is calculated using equation (37) for continuous distributions [27]. The integral is replaced by a summation for discrete distributions. A true distance measure produces the same result no matter the direction of the distance measurement, $D(p\|q)=D(q\|p)$. In this regard the KL distance is not a true distance measure since this property is not adhered to. It is thus more technically correct to refer to the KL distance as the KL divergence.

$$D(p\|q)=\int_x p(x)\log\left(\frac{p(x)}{q(x)}\right)dx \tag{37}$$

Kullback himself prefers the KL distance to be referred to as minimum discrimination information (MDI) [28]. In literature this preference is not used as many papers still refer to MDI as the KL distance. In keeping with this nomenclature and for the reason described previously, this report will use the name KL divergence. The disadvantage of this method is that it does not provide information on the variability of the results from test to test. The chi-squared method remedies this shortcoming.

The chi-squared (χ^2) test produces a single figure for how well the generated distribution matches the desired distribution. The χ^2 goodness of fit test, tests if the deviation from an expected distribution is due sampling error or if there is an actual difference. Sampling error is expected from the nature of statistics while a real difference would be an indication of a poorly generated distribution. The null hypotheses H_0 states that there is no significant difference between the generated distribution and the proposed distribution while the alternate hypotheses H_1 states that they are significantly different. A 95% or even a 99% confidence interval (CI) is commonly used, which is a measure of the certainty with which the alternate hypotheses can be rejected. It is important to note that the χ^2 test measures whether two distributions are significantly different, not if they are significantly similar, and that the lack of the fact that they are significantly different does

not mean they are thus significantly similar. This can lead to a scenario where the proposed distribution as a model can be rejected with 95% confidence, but cannot be rejected with 99% confidence. Thus lower percentage values in a χ^2 test are better. The main drawback of the χ^2 test is that the results are dependent on the size of the bins of the histogram under test. A general rule of thumb is to have bin sizes of 0.3 times the standard deviation. The most important factor is to have at least 5 samples in a bin. This means bins at the tails of the distribution might have to be enlarged and combined to ensure that this lower limit on the number of samples is adhered to [29], [30]. Equation (38) is used to calculate the chi-squared value χ^2. O_k is the observed value for a bin, E_k is the expected value (from the analytical function in this case) and n is the total number of bins.

$$\chi^2 = \sum_{k=1}^{n} \frac{(O_k - E_k)^2}{E_k} \tag{38}$$

Skew and kurtosis can also be used to measure the quality of the generated Gaussian random variable. Skew is the normalised third moment of the distribution and kurtosis is the normalised fourth moment of the distribution and is calculated using equations (39) and (40) respectively. The excess kurtosis is defined as the kurtosis minus 3, given in equation (41).

$$Skew = \frac{\mu_3}{\sigma^3} = \frac{\frac{1}{N}\sum_{n=1}^{N}(x_n - \bar{x})^3}{\left(\frac{1}{N}\sum_{n=1}^{N}(x_n - \bar{x})^2\right)^{3/2}} \tag{39}$$

$$Kurtosis = \frac{\mu_4}{\sigma^4} = \frac{\frac{1}{N}\sum_{n=1}^{N}(x_n - \bar{x})^4}{\left(\frac{1}{N}\sum_{n=1}^{N}(x_n - \bar{x})^2\right)^{2}} \tag{40}$$

$$Excess\ Kurtosis = \frac{\mu_4}{\sigma^4} - 3 \tag{41}$$

With respect to these equations N is the number of samples, x_n is the n-th sample, and $\bar{x}$ is the sample mean. The skew of the distribution is a figure that gives an indication to how symmetric the distribution is. A negative value for the skew translates to a distribution that is skewed to the left while a positive value translates to a distribution that is skewed to the right. A distribution that is skewed to the left has a longer tail on that side than the other. Kurtosis is a measure that gives an

indication of the shape of the distribution. A perfect Gaussian distribution has a kurtosis of 3. When the kurtosis is larger than this value, it indicates a distribution that is flatter and wider. A smaller kurtosis indicates the that a distribution is more peaked, or sharp [31].

For clutter data analysis higher order moments are often used to measure the accuracy of the distribution. The higher the order of the moment the more it is influenced by the values in the tails of the distribution [32]. Since the largest mismatch between a distribution and actual clutter data is in the tails, this property makes the higher order moments ideal for analysing how well a distribution matches the measured data. The moments up to order 6 are typically used for this analysis, since the calculated moments of order 3 and less can still be matched fairly easily to many different distributions (allowing for a small error) [33]. Not all error measurements occur on data of a statistical nature.

The error between two non statistical functions can be calculated by two different distance methods. The first method is the relative error, and the second is the absolute error. The relative error is calculated by subtracting the analytical PDF from the measured PDF, and then dividing by the analytical PDF. The percentage error is the relative error multiplied by 100. The absolute error is calculated similarly, except the step of dividing by the analytical PDF is left out. The error calculation methods are given in equation (42) for the relative error and equation (43) for the absolute error.

$$\epsilon(x)_{relative} = \frac{P_m(x) - P_a(x)}{P_a(x)} \tag{42}$$

$$\epsilon(x)_{absolute} = P_m(x) - P_a(x) \tag{43}$$

Using the measures of relative error, absolute error, χ^2, KL divergence value, skew, kurtosis and the higher order moments, the accuracy of the random number generator can be thoroughly investigated.

3.4 Gaussian Random Number Generator

The proposed generation methods for the Rayleigh, Log-normal and Weibull distributions for application in hardware clutter simulation all depend on a Gaussian random variable as will be seen in the sections which follow. First results will be obtained from a trusted Gaussian random number generator by means of simulation, which will serve as a baseline against which comparisons can be made.

3.4.1 Experiment 1: Baseline uncorrelated Gaussian random number generator

The first experiment is the analysis of a trusted Gaussian random number generator. As few as possible approximations are made which allows this experiment to serve as a benchmark for the ones to follow. The absolute error, relative error, χ^2 error, skew, kurtosis and higher order moments of the distribution are measured. The bins with less than 5 samples have been combined automatically in the analysis code written to generate these results. A Gaussian variable with a mean of 0 and a standard deviation of 1 was generated. For the simulation one million samples were generated.

Figure 23 is a plot of the generated Gaussian samples with an overlay of the analytical function of the distribution. The generated Gaussian samples plot has been normalised so that the area underneath it matches the area of the analytical function, by dividing the bin counts by the total number of samples generated. The area underneath the analytical function is equal to one by definition of the PDF. Visually one can observe that the generated distribution and the analytical function are a close match. There is however no measure of how well they match, which is required if a scientific comparison is to be made to other results. A visual measure like this is not objective or quantifiable, and tends to be very noisy. It does not have a high enough degree of accuracy and will produce different results from experiment to experiment.

Figure 24 shows a plot of the relative error and the absolute error. It is visible from the graphs that the absolute error and relative error methods of error measurement both have their weaknesses. The relative error becomes larger towards the tails of the distribution, while the absolute error becomes smaller towards the tails of the distribution. Close to the mean the relative error and absolute error performance is roughly the same.

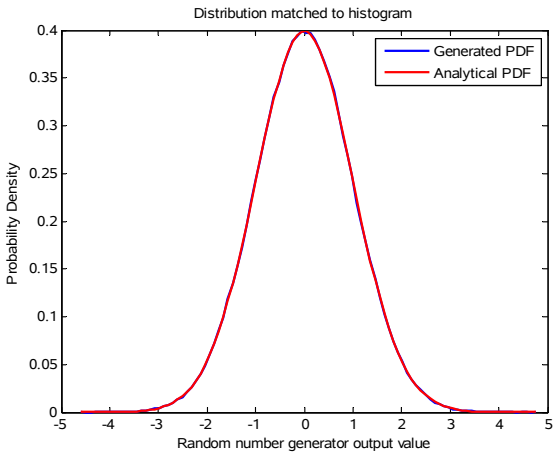

Figure 23: Normalised histogram of the simulated Gaussian random variable.

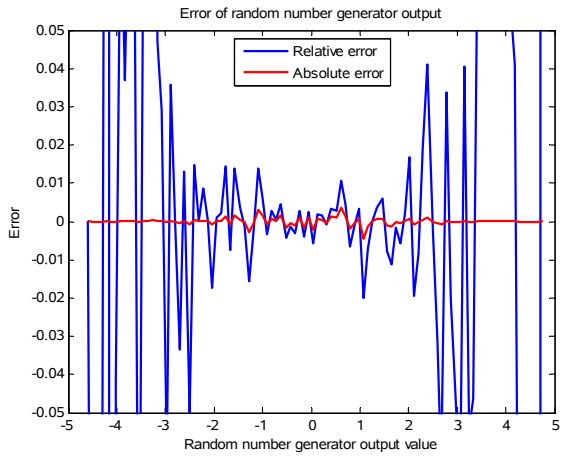

Figure 24: Comparison of the relative error (in blue) and the absolute error (in red).

These methods of evaluating the performance of the distribution close to the tails are insufficient, because they lack accuracy. These types of error measurement are also not statistically valid methods for evaluating the error between the generated and desired PDF. The Kullback-Liebler (KL) divergence test is a statistically valid method [27], and was used as a distance method to calculate the error between the generated and desired PDF. Performing the KL divergence test on the uncorrelated Gaussian random samples averaged over 30 experiments resulted in a value of 0.00010147. Although this value by itself does not give a clear indication of how well the samples generated matched the analytical PDF, this will become clear as results for the experiments that follow are obtained. It is important to note that the KL divergence test is sensitive to the number of bins and the number of samples used to calculate its value. It is thus important to fix these values for the experiments that follow to ensure the comparison is fair.

Figure 25 is a plot of the frequency content of the random numbers. The figure shows the frequency spectrum of the Gaussian random variable normalised to sampling frequency with the magnitudes taken relative to the largest frequency component. For all practical purposes the spectrum is white. There are nulls in the frequency spectrum at 0 Hz and at multiples of half the sampling frequency. The cause for these nulls are due to the fact that the frequency spectrum for a real signal is symmetrical. This effect is limited to the Gaussian distribution when looking through all the results presented, because this is the only real signal analysed in this dissertation. The Gaussian distribution is also the only symmetric distribution analysed. The frequency spectrum of the signal is obtained by splitting the time data into a number of sections having the same length. A fast Fourier transform (FFT) is then performed separately on each of these sections of time data. The absolute value of all of the separate FFT data is then averaged to obtain a smooth graph. The final FFT is then adjusted by placing the largest magnitude at 0 dB. This simplifies measurement of the -3 dB point on the graphs.

Figure 26 shows the autocorrelation plot of the sequence of random numbers. The autocorrelation is completely correlated and maximum at time equal to zero, but decorrelated directly before and directly after that point. Three data points are highlighted on the graph to show that the decorrelation time is within a single sample. This is an indication that there is no correlation between consecutive generated random numbers. This meets the expectations for the baseline experiment.

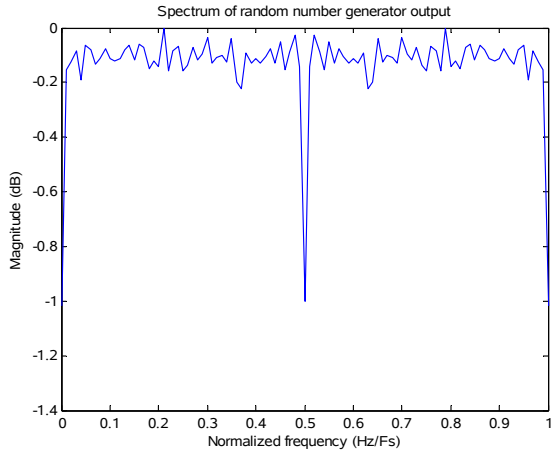

Figure 25: Frequency spectrum of the Gaussian random variable.

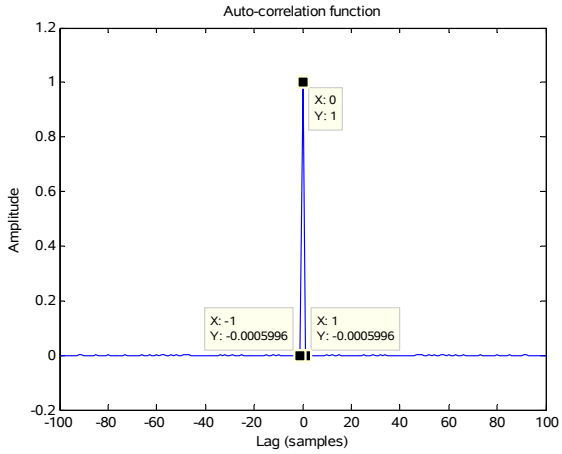

Figure 26: The autocorrelation function of the random variable.

Calculating the moments of the generated samples provides additional insight into the shape of the distribution. The desired theoretical values and measured values for the mean, standard deviation, variance, skew and kurtosis are shown in Table 4. The table also includes the results from the χ^2 test with the calculated value for a 95% confidence interval listed in the "expected" row, and the computed result shown in the "measured" row. For the χ^2 hypotheses test values the Sankaran closed form analytical approximation for the cumulative distribution function was used [34], [35]. This approximation produces values that are slightly smaller than the χ^2 pre-calculated tables in text books. The number of bins in the histogram k_T was set to 100, and since the degrees of freedom v is $v = k_T - p - 1$ where p is the number of parameters of the hypothesised distribution (2 for Gaussian distribution), the value for v is 97. As discussed in Section 3.2 the number of samples in a bin must be at least 5. To meet this requirement the values in the tails of the distribution were combined as a single bin when they were larger than 4 times the standard deviation. As a result the value for v was usually in the range of 80 to 90.

Table 4: Measured values over a million random numbers generated according to the Gaussian distribution.

	Mean	Standard deviation	Variance	Skew	Kurtosis	χ^2 $\alpha = 0.05$
Analytical	0	1	1	0	3	105 ($v = 83$)
Measured	0.0011	0.9996	0.9991	-0.0021	2.9998	76.4 (H_0 accepted)

The moments of order 1 to 6 were calculated analytically and compared to the same order moments measured from the generated data. The comparison was done by calculating the percentage error between the measured and calculated moments. For the Gaussian distribution all odd moments are

equal to zero, which results in an infinite percentage error. The odd numbered higher order moments have been tabulated as not applicable (NA) in Table 5 along with the percentage error for the even numbered higher order moments.

Table 5: Percentage error of the higher order moments for the uncorrelated Gaussian distribution.

Moment order	1st	2nd	3rd	4th	5th	6th
Percentage error	NA	0.1287%	NA	0.0801%	NA	0.2078%

Figure 27 shows the calculated theoretical moments plotted against the measured moments from the Gaussian random number generator. All odd order theoretical moments are equal to zero and thus do not appear on a log scale. Both the percentage error and the figure confirms that there is an excellent match between the generated data and a Gaussian distribution.

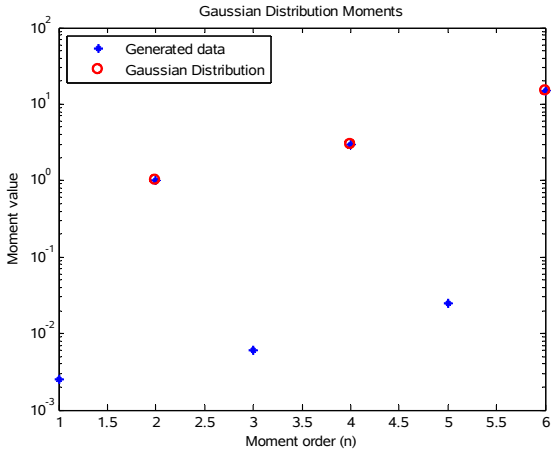

Figure 27: Measured moments from the generated Gaussian distribution.

Confidence intervals were used to obtain a statistical measure of the variation of the results and their significance. The calculations were performed for a confidence level of 99%. A double sided confidence was calculated for values that could lie within a range (i.e. values that need to be limited from both the upper and lower side), which for this case was every value except the χ^2 test result. The χ^2 test result only had an upper confidence bound calculated because the calculated χ^2 must be less than the theoretical limit, and is thus only bound from the upper side. The experiment was repeated 30 times to obtain the confidence intervals as shown in Table 6. Confidence intervals of 99% were calculated. It can be clearly seen that there is a large amount of certainty in the values of the mean, standard deviation, variance, skew and kurtosis since the confidence intervals are small and thus very close to the means of the results. The χ^2 test result is a bit less certain. Out of the 30 experiments the χ^2 test could not reject the null hypotheses H_0 28 times, which means there were 2 cases where it was rejected.

Table 6: Results of the experiment as tabulated in Table 4 repeated 30 times.

	Mean	Standard deviation	Variance	Skew	Kurtosis	χ^2	χ^2 lim.
Lower CI	-0.00057	0.99963	0.99925	-0.00068	2.99615	N/A	101.769
Mean	-0.00011	1.00002	1.00005	0.00052	2.99853	82.4731	103.304
Upper CI	0.00035	1.00042	1.00084	0.00171	3.00091	88.9083	104.841

If only one experiment was performed, and thus a single result obtained, it would be impossible to know if it this was due to chance or because the random number generator is failing to generate an accurate enough distribution. The confidence intervals on the χ^2 test boundaries are very small, and this is expected since they are related to the number of degrees of freedom. The number of degrees of freedom only vary due to the combination of bins in the tails of the distribution. The upper values of the calculated χ^2 test result are well below lower limit of the χ^2 test boundary, which confirms that the random number generator is an accurate number generator for the Gaussian distribution.

3.4.2 Experiment 2: Weakly correlated Gaussian random number generator

This experiment tests the output after correlation has been added to the Gaussian random number sequence, without any other modification to the algorithm. Correlation is added by filtering the output of an uncorrelated Gaussian random number generator as discussed by the filter approach outlined in [12] in Section 3.2. The results of the filter bandwidth have been normalised to the sampling frequency of that filter, to make them independent of the sampling frequency.

This experiment was performed with a normalised bandwidth of $0.1\,Fs$ where Fs is the sampling frequency of the filter. Figure 28 shows that the generated random number generator follows the PDF of the analytical Gaussian distribution closely. The KL divergence value obtained when averaging over 30 experiments was 0.00010383. This value is 2.33% greater than that of the uncorrelated case obtained for the previous experiment. Because a smaller value is better, this is an indication that when correlation is added to the random sequence the performance could be degraded. A more thorough investigation of this result is required before a valid conclusion can be made.

The frequency spectrum plot in Figure 29 shows that the -3 dB bandwidth is close to the required $0.1\,Fs$. The magnitudes were taken relative to the largest frequency component. There is a slight notch in the frequency spectrum at $0\,Fs$, $1\,Fs$ and at $0.5\,Fs$. The size of the notches in the spectrum are around 1 dB, the same as for the uncorrelated case in Section 3.4.1.

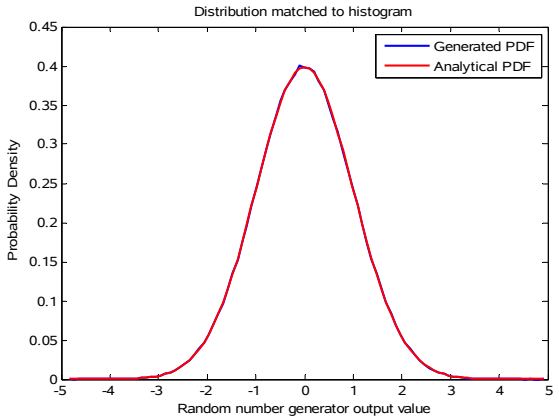

Figure 28: Normalised histogram of the correlated Gaussian random variable.

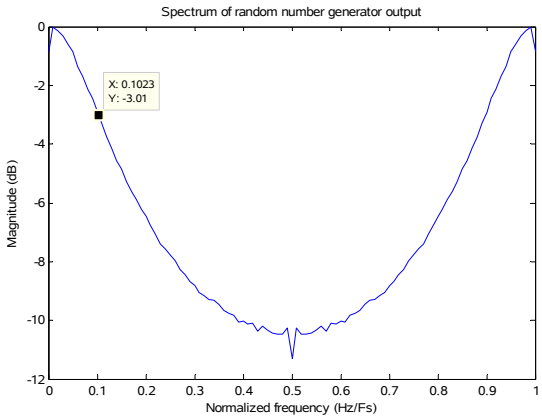

Figure 29: Frequency spectrum of the correlated Gaussian random variable.

The autocorrelation function in Figure 30 shows that there is correlation from sample to sample. The correlation times were measured as 1.2 samples for $\tau_{1/2}$ and 1.7 samples for $\tau_{1/e}$ as shown with the data points on the plot. The autocorrelation function is a maximum at a lag of 0 samples which is the theoretical position of its maximum. The sequence was completely decorrelated after approximately 10 samples.

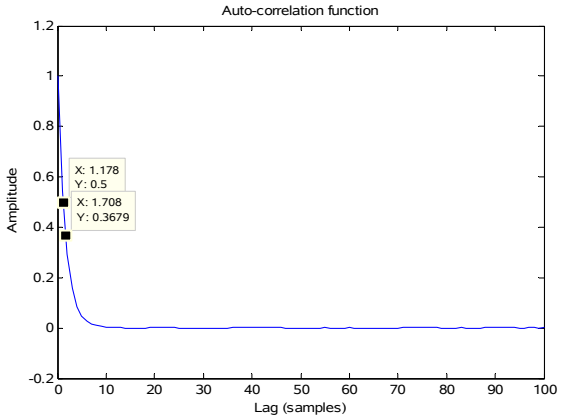

Figure 30: The autocorrelation function of the random variable.

The confidence interval results shown in Table 7 were the result of the above experiment repeated 30 times. The confidence intervals were calculated at 99%. The confidence intervals for the mean, standard deviation, variance, skew and kurtosis is still very close to the mean values as they were with the uncorrelated experiment in 3.4.1, which indicates a high degree of certainty in those values. Even though there is still a high degree of certainty the confidence intervals are further apart than they were previously, an indication that the certainty is less but negligibly so. The χ^2 test limit is not much different from the previous experiment and the average χ^2 test value is slightly less.

Out of the 30 repeated experiments the χ^2 test could not reject the null hypotheses H_0 29 times. The higher order moments were recorded for the correlated Gaussian experiment with a bandwidth of 0.1 Fs. The higher order moments are tabulated in Table 8.

Table 7: Results of the simulation of correlated Gaussian numbers with a bandwidth of 0.1 Fs.

	Mean	Standard deviation	Variance	Skew	Kurtosis	χ^2	χ^2 lim.
Lower CI	-0.0011	0.9996	0.9991	-0.0021	2.9970	N/A	101.0035
Mean	-0.00029	1.0001	1.0001	-0.00049	2.9998	80.9965	103.1890
Upper CI	0.00051	1.0006	1.0011	0.0011	3.0027	87.2587	105.3745

Table 8: Percentage error of the higher order moments for the weakly correlated Gaussian distribution.

Moment order	1st	2nd	3rd	4th	5th	6th
Percentage error	NA	-0.1947%	NA	-0.3511%	NA	-0.1384%

The percentage error of the higher order moments for the correlated Gaussian is larger than that of the uncorrelated case, but the error is still very small. The percentage error thus confirms that there is still a satisfactory match between the generated data and a Gaussian distribution, even if the samples of the Gaussian distribution are correlated by filtering.

3.4.3 Experiment 3: Strongly correlated Gaussian random number generator

The second experiment was performed with a bandwidth of $0.01\,Fs$. Figure 31 shows that the generated random number generator follows the PDF of the analytical Gaussian distribution closely.

The KL divergence value obtained when averaging over 30 experiments was 0.00013393. This value is 31.99% greater that the uncorrelated case. This error is also approximately 10 times greater than the weakly correlated case in the previous experiment. The addition of correlation to the random samples thus not only decreases the accuracy of the PDF, but this decrease in accuracy is worse for stronger correlated samples.

The frequency spectrum plot in Figure 32 shows that the -3 dB bandwidth is not at the required $0.01\,Fs$ but instead closer to $0.015\,Fs$ which is a 50% error. The spectrum is normalised with regards to the sampling frequency, and the magnitudes taken relative to the largest frequency component. The measured -3 dB mean value was 0.0146, with a 99% confidence interval of ±0.0001. The autocorrelation function in Figure 33 shows that there is correlation. The autocorrelation function is a maximum at a lag of 0 samples. The correlation times were measured as 11 samples for $\tau_{1/2}$ and 15.9 samples for $\tau_{1/e}$. The signal was completely decorrelated after approximately 100 samples.

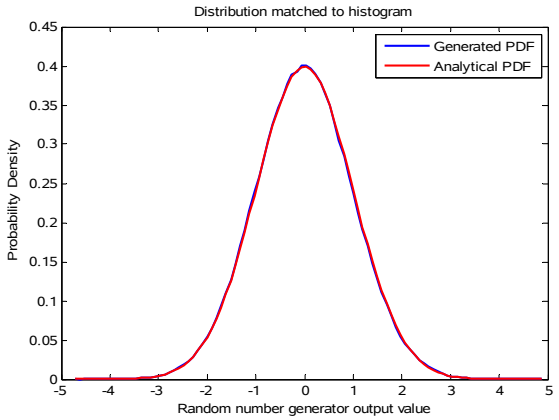

Figure 31: Normalised histogram of the correlated Gaussian random variable.

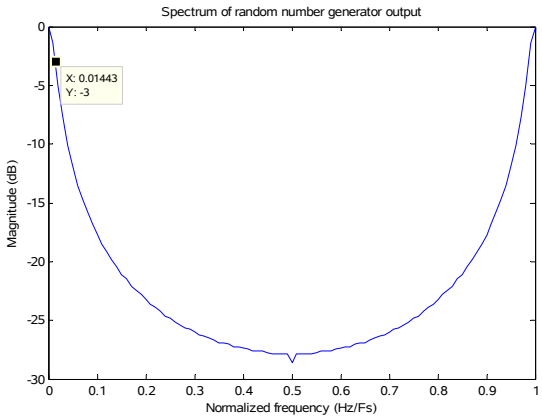

Figure 32: Frequency spectrum of the correlated Gaussian random variable.

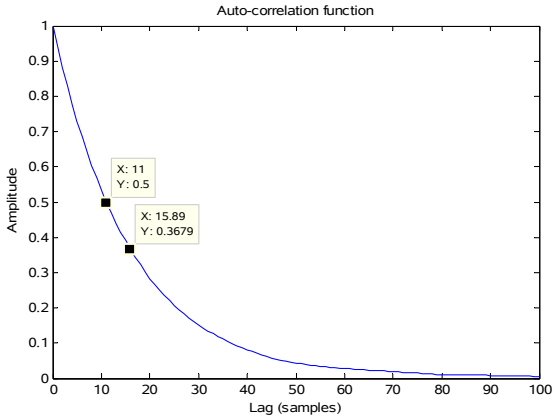

Figure 33: The autocorrelation function of the random variable.

The confidence interval results shown in Table 9 were the result of the above experiment repeated 30 times. The confidence intervals were calculated at 99%. The confidence intervals for the mean, standard deviation, variance, skew and kurtosis is still very close to the mean values as they were with the previous correlated experiment in 3.4.2 and the uncorrelated experiment in 3.4.1, which indicates a high degree of certainty in those values. Even though there is still a high degree of certainty the confidence intervals are further apart than they were for both the previous correlated experiment in 3.4.2 and the uncorrelated experiment in 3.4.1. There is thus a decrease in the certainty of the results. The χ^2 test limit is not much different from the previous experiments. The average χ^2 test value was very high at close to double the required value. Out of the 30 repeated experiments the χ^2 test could not reject the null hypotheses H_0 only twice. This is a dramatic decrease in accuracy of the generated distribution.

This failure of the χ^2 test is a surprise since there is no indication from the moments that the generated distribution is significantly different from the previously generated ones. The error curves do hint at a problem though with relatively large errors. Even though the error curve is unbalanced

the moments do not steer one to believe that there is anything wrong with the shape of the distribution. The fact that the plot of the average error over 30 experiments is not symmetric is also of a concern which means that this underlying error is not random but a bias present in every sequence generated. The higher order moments were recorded for the correlated Gaussian experiment with a bandwidth of 0.01 Fs . The higher order moments are tabulated in Table 10.

Table 9: Results of the simulation of correlated Gaussian numbers with a normalised bandwidth of 0.01.

	Mean	Standard deviation	Variance	Skew	Kurtosis	χ^2	χ^2 lim.
Lower CI	-0.0040	0.9985	0.9970	-0.0041	2.9971	N/A	102.4282
Mean	-0.0011	1.0000	1.0000	0.00002	3.0047	183.4319	104.8062
Upper CI	0.0018	1.0014	1.0029	0.0041	3.0123	209.2534	107.1843

Table 10: Percentage error of the higher order moments for the strongly correlated Gaussian distribution.

Moment order	1st	2nd	3rd	4th	5th	6th
Percentage error	NA	-0.9129%	NA	-1.7173%	NA	-1.4180%

The percentage error for the 0.01 Fs correlated Gaussian is larger than that of the 0.1 Fs correlated Gaussian and the uncorrelated case, but the error is still very small. Figure 34 shows that there is no visible difference between the measured and calculated theoretical moments, even for the 6th moment with the largest percentage error. All odd order theoretical moments are equal to zero and do not appear on a log scale.

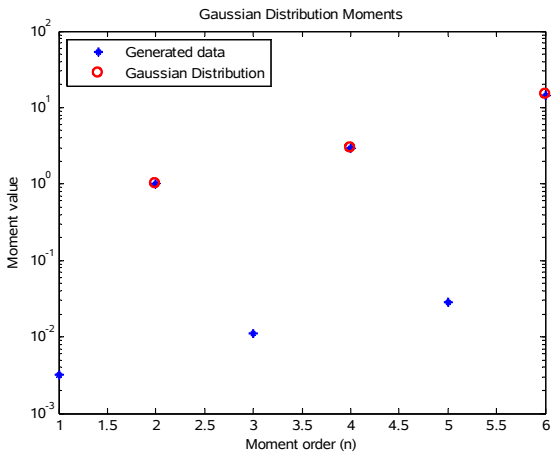

Figure 34: Measured moments from the generated correlated Gaussian distribution.

The percentage error thus confirms that there is still a satisfactory match between the generated data and a Gaussian distribution, even if the samples of the Gaussian distribution are correlated by filtering.

3.4.4 Experiment 4: Trends with varying degrees of correlation for the Gaussian random number generator

From the previous two experiments in sections 3.4.2 and 3.4.3, it was observed that there is a degradation in performance of the method used to generate the correlated Gaussian random samples. This degradation is visible in the accuracy of the bandwidth, the accuracy of the PDF (measured with the KL divergence), and the accuracy of the higher order moments. This experiment investigates the trend of that degradation in performance, by varying how strongly correlated the generated samples are. The desired bandwidth is stepped from 0.01 Fs to 0.25 Fs, with all other variables kept constant.

From the curves in Figures 35 and 36 it can be seen that the expected -3 dB value as set for the correlated Gaussian random number generator differs from the actual -3 dB value as the expected -3 dB bandwidth chosen value becomes smaller. The 99% confidence intervals of the measured -3 dB point are almost exactly the same as the mean for the measured -3 dB point. This indicates that the deviation from the expected value is not due to random chance, but due to some yet unknown process. The error for the smallest normalised bandwidth of 0.01 is almost 50% which is significant, but it decreases rapidly for larger bandwidths.

The data from 100 repeated χ^2 tests at 25 different normalised bandwidths, shows the emergence of a clear trend.

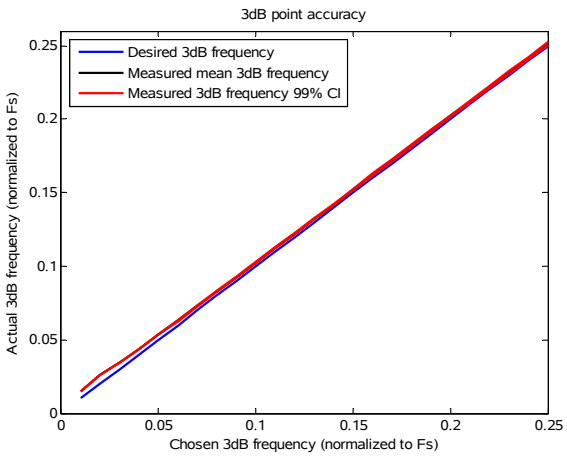

Figure 35: Measured -3 dB bandwidth of the correlated Gaussian random number.

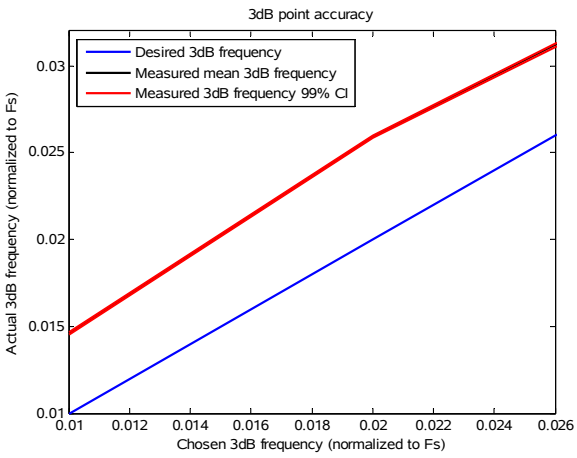

Figure 36: A zoomed in version of the graph of Figure 35 showing narrow bandwidth settings.

Figure 37 shows the average value of the outcome of the 100 repeated χ^2 tests at each bandwidth setting. The mean outcome is shown in black, the 99% confidence interval in red and the average value for the test rejection threshold in blue. The average value of the mean outcome is still less than the threshold at about 0.025, but the 99% confidence upper bound only crosses this line at about 0.05 Fs. From Figure 38 it can be seen that the χ^2 test success rate is acceptable for larger

bandwidths, and almost constant as low as 0.075 Fs, where after it starts to decline sharply. The χ^2 test thus fails to reject the Rayleigh PDF with 99% confidence between bandwidth values of 0.05 Fs and 0.25 Fs.

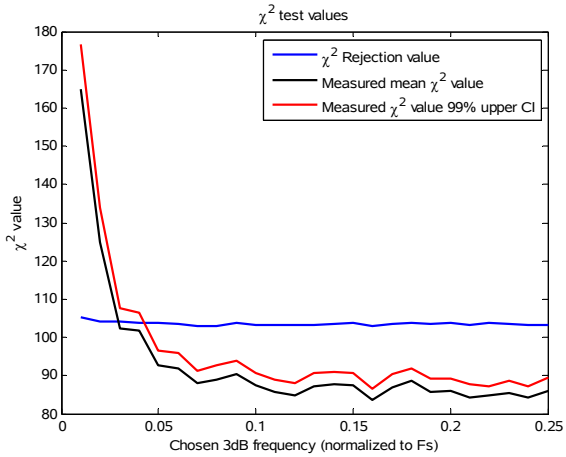

Figure 37: Plot of the outcome from a chi squared test.

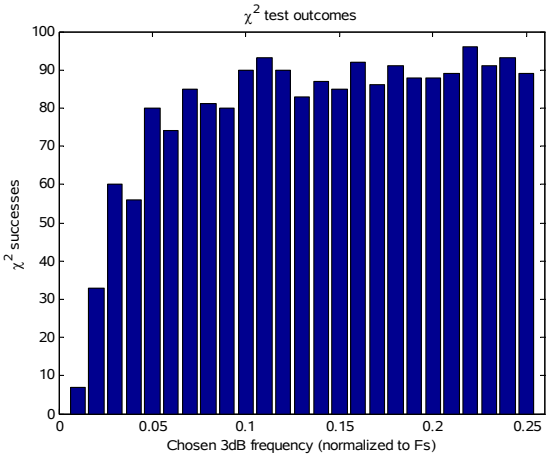

Figure 38: Number of chi squared tests that were unable to reject the PDF.

Figure 39 shows the KL divergence value with varying degrees of correlation. There is a trend of worsening performance as the correlation in the samples is increased. This confirms the trend observed in the previous two experiments in this section. There appears to be no increase in the variance of the results obtained, although the largest outlier worsens as indicated by the red dashed line in the plot. Figure 40 shows the number of samples in the two tails of the Gaussian distribution, summed to give a value out of 100. There is no clear trend, and the number of samples in the tails appears to be constant irrespective of the normalised bandwidth setting. There is however an increase in variability of the result as seen by the variance shown on the plot by the blue lines, and the increase in maximum and minimum outliers shown on the plot by the red lines respectively. So even though the mean number of samples in the tail is accurate, the result from a single experiment starts to make a larger and larger error as the correlation is increased.

Figure 41 shows the results from of the percentage error between the measured higher order moments and the calculated ones. All odd order moments produce an infinite percentage error and is not shown. The trend as observed for the χ^2 test result in Figure 37, and the KL divergence value in Figure 39 does not emerge for the higher order moments. All three the measured moments have the same trend with an error varying around zero percent irrespective of the bandwidth setting. The larger the order of the moment the larger the percentage error. To investigate if there are possibly underlying trends in Figure 41 that is not visible because of under sampling in the frequency steps chosen, and to investigate if there is a possible trend below the 0.01 Fs setting, Figure 42 was generated with steps of 0.001 Fs from 0.001 Fs to 0.05 Fs. As can be seen from this plot there is

still no trend in the mean with the errors varying around 0 percent. The only visible trend is in the varying size of the errors which increases for smaller bandwidth settings, which was also observed with the KL divergence value.

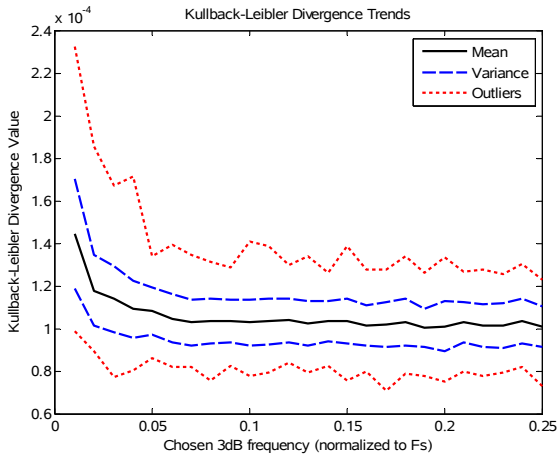

Figure 39: Kullback-Leibler divergence value.

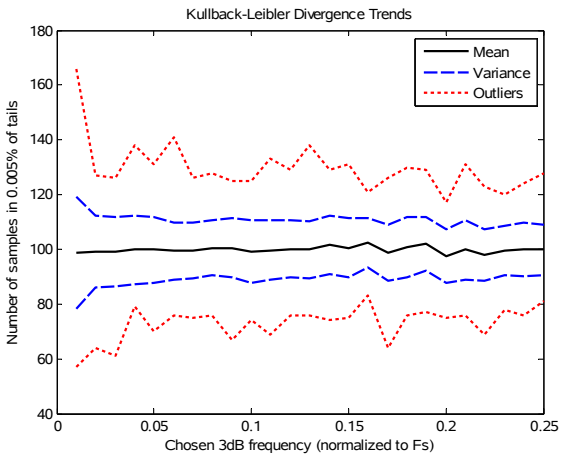

Figure 40: Number of samples in 0.005% of the tails on both sides combined.

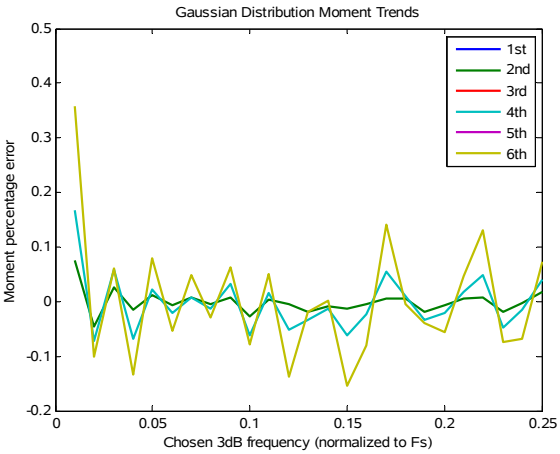

Figure 41: Percentage error of the measured moments from the generated correlated Gaussian distribution.

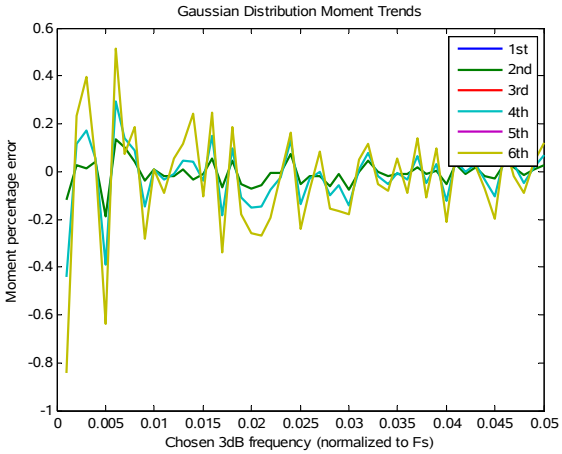

Figure 42: Percentage error of the measured moments from the generated correlated Gaussian distribution for a smaller range of frequencies.

These tests indicate that there is a deviation from the ideal Gaussian PDF when the normalised bandwidths are set to small values. This deviation from the ideal PDF could also be responsible for the error in the -3 dB setting as discussed in the previous paragraph. The Rayleigh PDF is simulated from two Gaussian PDFs, and the open question is how much the Rayleigh PDF is influenced by the irregularities of the Gaussian.

3.4.5 The effect of filter ramp up

The trends observed in the results with varying bandwidth in Figure 35 and Figure 36 shows that the correlation filter is increasingly inaccurate at smaller bandwidths and less so at larger bandwidths. One possible cause for this would be filter ramp up. The ramp up time is the time it

takes for the filter to reach its steady state. During the start up of the filter, the output statistics can be considered as inaccurate, and these samples are often discarded. The narrower the bandwidth of a filter, the longer it takes to reach its steady state. This can easily be seen from the impulse response of the filter. An investigation to determine if filter ramp up was the culprit was performed.

Removing a predetermined fixed number of initial samples will not give accurate results, since the ramp up time increases as the filter bandwidth decreases. For this experiment the point where the samples are considered to be accurate enough, was based on the step response input. The step response for each bandwidth setting was calculated, and the slope between successive samples within the step response recorded. When the slope reached 1×10^{-6} the sample number was saved and multiplied by a thousand. This value was used as the starting point of the random clutter sequence for the test where the ramp up was removed. One million samples were recorded after this point as the random number sequence under test. For the sequence where the ramp up was not removed the first million samples were used as discussed in the previous experiments.

Figure 43 and Figure 44 show the result after ramp up was removed. From Figure 43 it can be seen that there is no visible difference between the bandwidths measured where the ramp up was not removed, and the samples where the ramp up was removed. Figure 44 shows a zoomed in version of this result, where there is a small difference between the two results. In this figure the result where the ramp up has been removed is slightly better than the result where the ramp up has not been removed. At other points in the figure the "with ramp up" trace is below the "without ramp up" trace, so the difference can be attributed to sample variation, and not the increased performance of one method compared to the other.

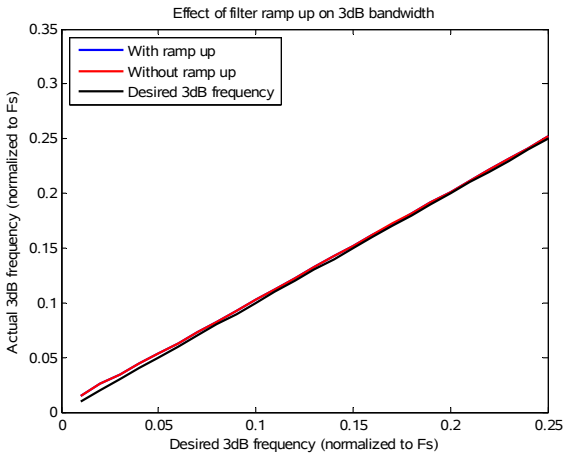

Figure 43: Comparison between discarding samples and not discarding samples.

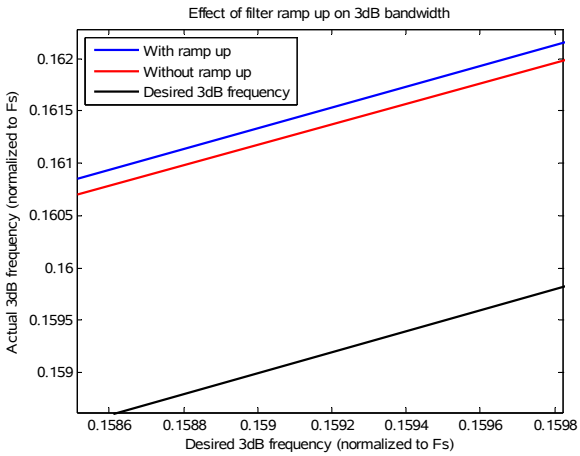

Figure 44: A zoomed in version of Figure 43 to highlight a difference between the two approaches.

Figure 45 shows the relative error between the desired and measured -3 dB bandwidth with and without filter ramp up. Figure 46 shows absolute error between the desired and measured -3 dB bandwidth with and without filter ramp up. From these figures it is clear that there is no real difference in the measured normalised -3 dB bandwidth.

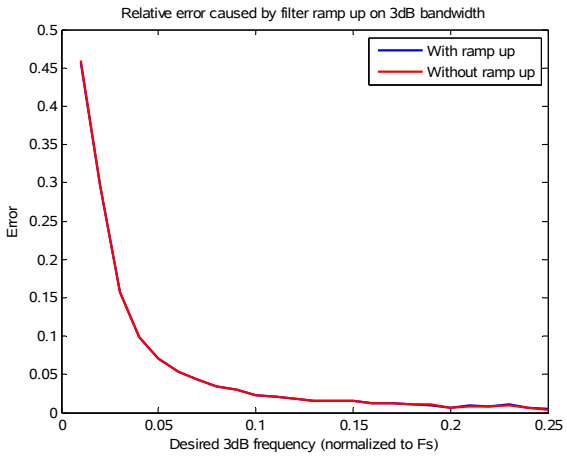

Figure 45: Relative error of the -3 dB bandwidth.

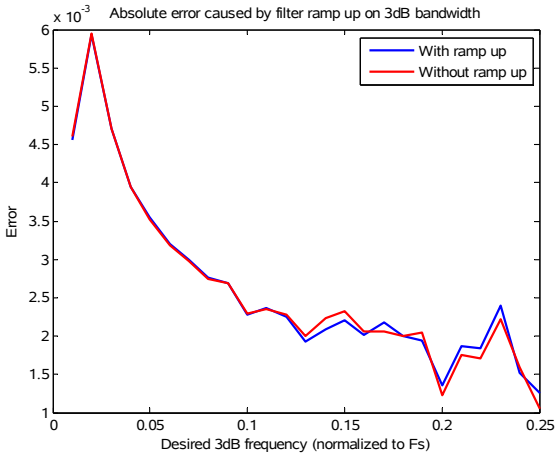

Figure 46: Absolute error of the -3 dB bandwidth.

The effect of filter ramp up on the results is thus negligible, and the cause for the error has to be found elsewhere. The results for the experiments that follow have not had the ramp up of the filter removed, as it was deemed an unnecessary waste of computation time.

3.5 Rayleigh clutter

The Rayleigh distribution described analytically in equation (10) from Section 2.3.3, repeated here as equation (44), is usually an acceptable description of clutter statistics when the range resolution of the radar is low.

$$P(x;\alpha)=\frac{x}{\alpha^2}e^{\frac{-x^2}{2\alpha^2}},\ x\geq 0 \tag{44}$$

A low radar resolution captures a large number of scatterers in a single range cell, which according to the central limit theorem, reduces the scattering behaviour to Gaussian statistics. The Rayleigh distribution is the direct cause of scatterers that behave according to the Gaussian distribution. By simply substituting independent Gaussian random variables into equation (22) to obtain the magnitude, a Rayleigh distribution results [36], [37]. The Rayleigh distribution is the least complex case to simulate, and does not require a memoryless non-linear transform (MNLT). To analyse the accuracy of the Rayleigh random sequence algorithm a baseline is created against which all the

analysis that follow will be measured. The PDF error, KL divergence value, frequency spectrum, phase distribution, autocorrelation function, χ^2 error and percentage error of the higher order moments of the resulting random sequences was analysed.

3.5.1 Experiment 1: Generation of the Rayleigh distribution using uncorrelated Gaussian random variables

This experiment serves as the baseline against which the experiments that follow related to the Rayleigh distribution will be compared. To accomplish this a well behaved Rayleigh random number sequence has been generated, with a value for α set to 1. There were one million uncorrelated samples generated for each Gaussian distribution on the I and Q channels, which resulted in one million uncorrelated Rayleigh distributed samples after obtaining the magnitude.

Figure 47 shows the analytical PDF plotted against the histogram of the generated samples from the Rayleigh random variable. There is a close fit between the analytical PDF and the measured samples. The phase of the random signal can be seen in Figure 48, and is uniformly distributed between $-\pi$ to π. A uniformly distributed phase is a requirement for clutter back-scatterer in a single range cell. A uniform phase results because of the relatively short wavelengths employed by radars. The small movement of scatterers between pulses can change the phase ϕ by more than 2π because the phase of the received signal and distance to the radar are related by $\phi = 4\pi R/\lambda$. The radar however can only measure phase in the 0 to 2π range. Combine this with the constructive and destructive summation of many scatterers in a range cell and it is clear how this distribution for the phase is obtained.

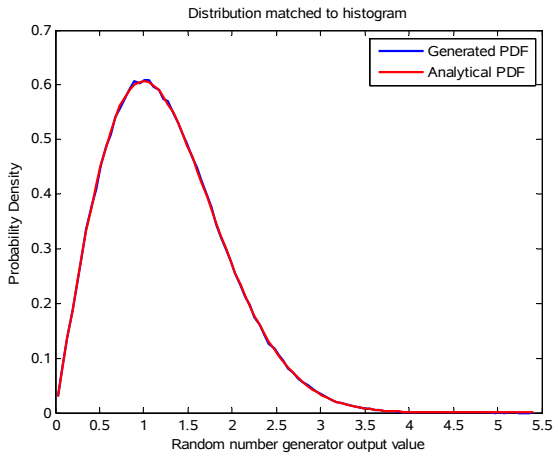

Figure 47: Normalised histogram of the uncorrelated Rayleigh random variable.

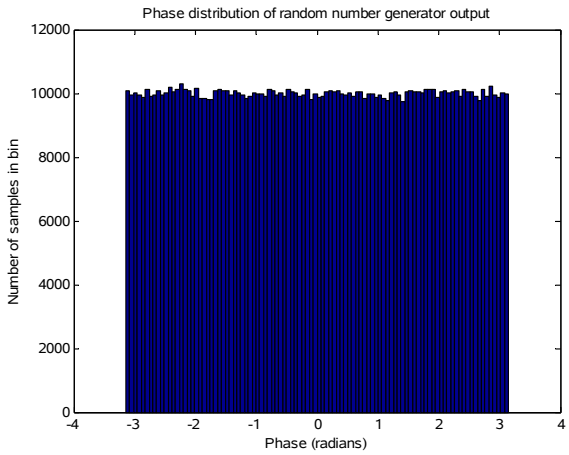

Figure 48: The phase of the uncorrelated Rayleigh random variable.

Performing the KL divergence test on the uncorrelated Rayleigh random samples averaged over 30 experiments resulted in a value of 0.00010035. This value is 1.104% less than that of the uncorrelated Gaussian distribution generated in Section 3.4.1. Because a smaller value is better, this is an indication that when two independent uncorrelated Gaussian distributed random numbers are added to produce a Rayleigh distribution the resulting Rayleigh distributed random sequence performs better than the individual Gaussian random sequences used. These results are very close to each other and not significantly different enough so that any form of concrete conclusion can be drawn from it. When running the 30 experiments a few times, the results for the Rayleigh random numbers would at times be much closer, and at other times worse than that of the underlying Gaussian random numbers. The variation in the results thus does not allow for a conclusion in either direction.

Figure 49 shows the frequency spectrum of the Rayleigh distributed number sequence. The frequency spectrum is normalised to sampling frequency with the magnitudes taken relative to the largest frequency component. The frequency spectrum is white, a result of an uncorrelated signal. Figure 50 shows the result from autocorrelation. Three data points are highlighted on the graph to show that the decorrelation time is within a single sample. The autocorrelation function is a maximum at a lag of 0 samples. The autocorrelation function is an impulse which serves as a confirmation that the sequence of random numbers is uncorrelated.

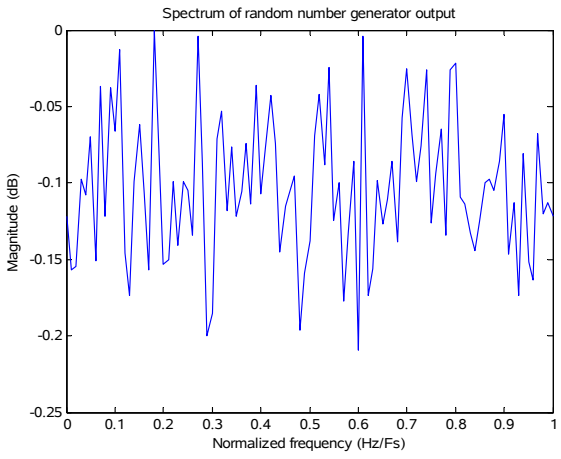

Figure 49: Frequency spectrum of the uncorrelated Rayleigh random variable.

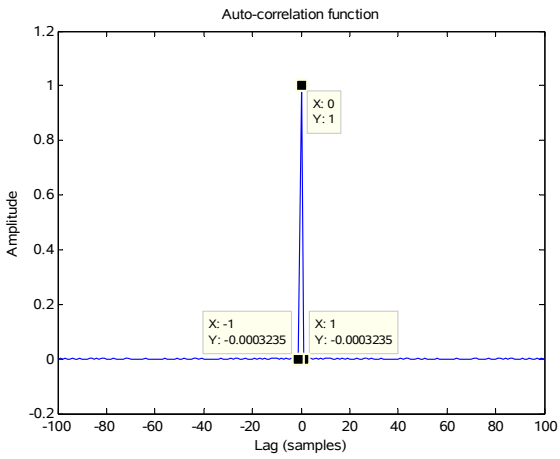

Figure 50: The autocorrelation function of the random variable.

The moments of order 1 to 6 were calculated and compared to the same order moments measured from the generated data. The comparison was done by calculating the percentage error between the measured and calculated moments. For the Rayleigh distribution the theoretical odd order moments are non zero because the distribution is not symmetrical. This is thus different from the Gaussian distribution. The higher order moments are tabulated in Table 11.

Table 11: Percentage error of the higher order moments for the uncorrelated Rayleigh distribution.

Moment order	1st	2nd	3rd	4th	5th	6th
Percentage error	0.0001%	0.0475 %	0.0984%	0.1387%	0.1663%	0.1812%

The percentage error for the generated Rayleigh moments are much smaller than the corresponding even order Gaussian moments, which supports the results of the KL divergence test that the Rayleigh distributed random numbers perform better than the Gaussian distributed random numbers they were generated from. Figure 51 shows the calculated theoretical moments plotted against the measured moments from the Rayleigh random number generator. Both the percentage error and the figure confirms that there is an excellent match between the generated data and a Rayleigh distribution.

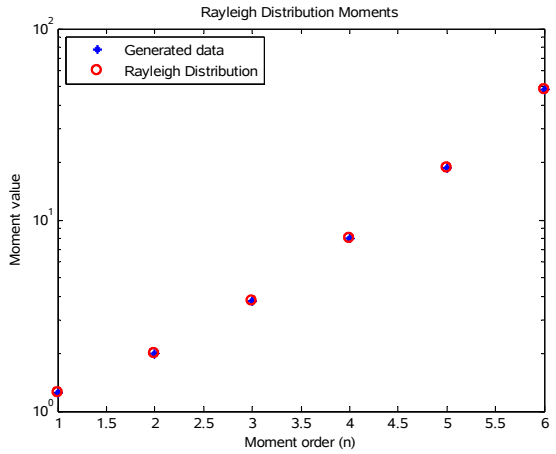

Figure 51: Measured moments from the generated Rayleigh distribution.

Repeating the experiment 30 times and performing a χ^2 test provides a measure of how well the PDF of the generated Rayleigh distribution fits to the required Rayleigh PDF. The χ^2 test was not rejected 28 times out of the 30 and on average the test result was not rejected with a confidence of more than 99%.

The baseline Rayleigh distribution analysed in this experiment performed well, and even outperformed the underlying Gaussian distributions from which it was constructed. The next step was to add a correlation to the random number sequence to analyse its effects on the performance of the distribution.

3.5.2 Experiment 2: Generation of the Rayleigh distribution using two independent time correlated Gaussian random variables

This experiment was performed with two independent sequences of Gaussian random numbers as described with equation (9). Each of the sequences were time correlated using the method discussed in Section 3.2 Although there was a correlation in the sequences with themselves there was no correlation between the two random number sequences, in other words the sequences were independent for I and Q. For this experiment a value of $\alpha = 1$ was used and the bandwidth of the signal was set to $0.01\,Fs$. One million correlated samples were generated for each of the Gaussian inputs to equation (9), to give one million correlated Rayleigh distributed samples.

Figure 52 shows the analytical PDF of the Rayleigh distribution compared to the correlated random samples generated as described in the previous paragraph. Visually there is a good match, except for some minor deviation at the maximum of the distribution. This deviation can comfortably be attributed to the stochastic nature of the generated samples. Figure 53 shows that the phase of the correlated Rayleigh sequence is uniform between $-\pi$ and π.

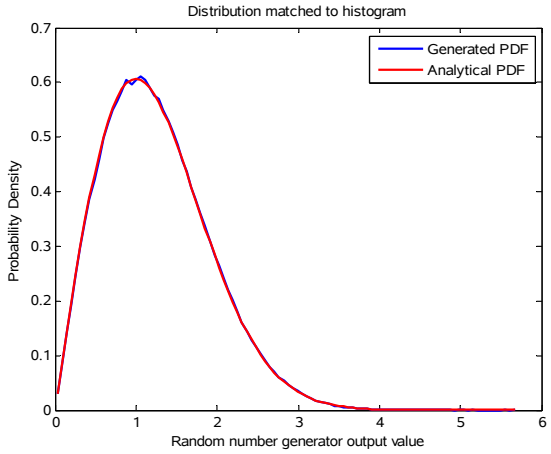

Figure 52: Normalised histogram of the time correlated Rayleigh random variable.

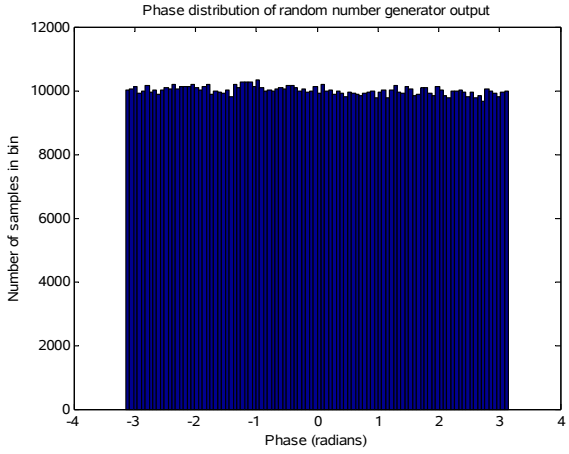

Figure 53: The phase of the correlated Rayleigh random variable.

The KL divergence test was used as a distance method to calculate the error between the generated and desired PDF. Performing the KL divergence test on the correlated Rayleigh random samples averaged over 30 experiments resulted in a value of 0.00011508, which is 14.07% better than the correlated Gaussian result. The KL divergence value of the correlated case is 14.78% worse than the uncorrelated case for the Rayleigh distribution, but the correlated Rayleigh distribution performs better than the correlated Gaussian with the same bandwidth setting, which is unexpected since the correlated Gaussian sequences which drives this experiment (Section 3.4.2), were used as input to the Rayleigh distribution. This is probably due to the fact that for the generation of the Rayleigh sequence, two Gaussian sequences are summed (as quadrature channels), and the error is thus only the average error between the two.

Figure 54 shows the actual bandwidth of the correlated Rayleigh sequence. The frequency spectrum is normalised to sampling frequency with the magnitudes taken relative to the largest frequency component. The normalised bandwidth was set to 0.01 Fs . The value measured at -3 dB was equal to 0.01183, an error of 0.00183, or 18.3%. The autocorrelation function in Figure 55 shows that the sequence is correlated. The correlation times were measured as 11.1 samples for $\tau_{1/2}$ and 16 samples for $\tau_{1/e}$. The autocorrelation function is a maximum at a lag of 0 samples. The signal was completely decorrelated after approximately 100 samples. These results are very similar to the results of the underlying correlated Gaussian sequence shown in Figure 33, and the discussion in Section 3.4.2. The higher order moments were recorded for the correlated Rayleigh experiment.

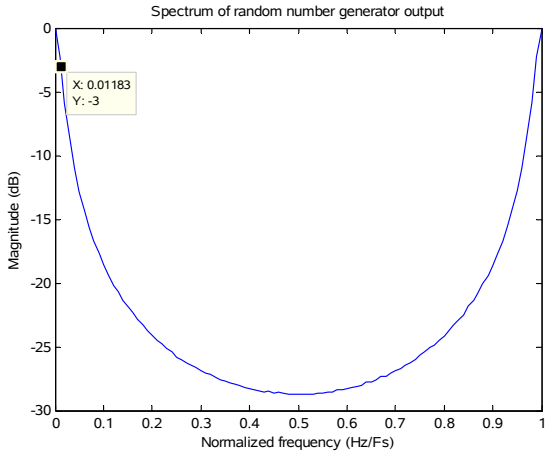

Figure 54: Frequency spectrum of the correlated Rayleigh random variable.

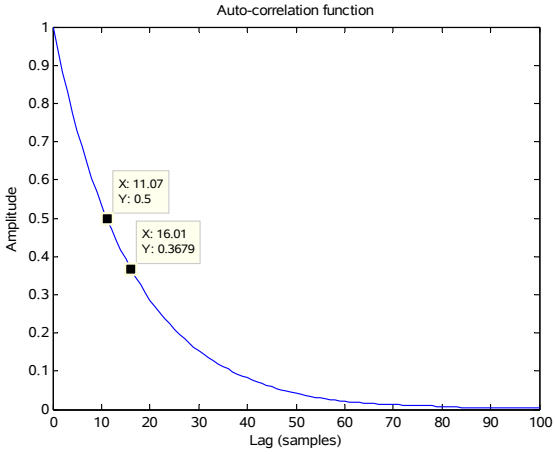

Figure 55: The autocorrelation function of the correlated random variable.

Figure 56 shows the calculated theoretical moments plotted against the measured moments from the Rayleigh random number generator. The higher order moments are tabulated in Table 12. The percentage error for the correlated Rayleigh distribution is larger than that of the uncorrelated case, but the error is still very small. The percentage error thus confirms that there is still an excellent match between the generated data and a Rayleigh distribution, even if the samples of the underlying Gaussian distributions are correlated by filtering.

Repeating the experiment 30 times and performing a χ^2 test provides a measure of how well the PDF of the generated correlated Rayleigh distribution fits the required Rayleigh PDF. The χ^2 test was not rejected 27 times out of the 30 and on average the test result was not rejected with a confidence of more than 99%.

Table 12: Percentage error of the higher order moments for the correlated Rayleigh distribution.

Moment order	1st	2nd	3rd	4th	5th	6th
Percentage error	-0.0052%	-0.0531%	-0.1615%	-0.3482%	-0.6286%	-1.0204%

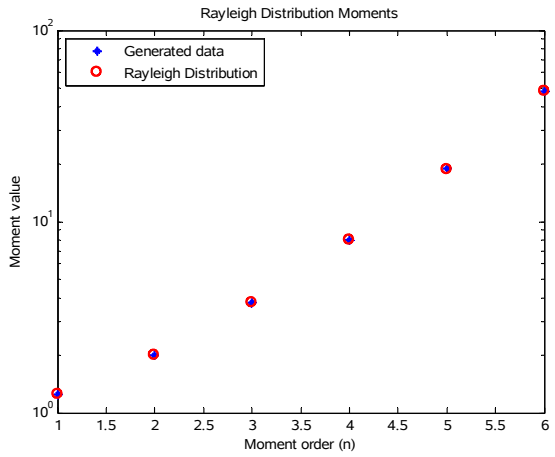

Figure 56: Measured moments from the generated correlated Rayleigh distribution.

The performance of the correlated Rayleigh distribution was only observed at a single bandwidth setting. The next experiment will investigate the performance across a wide range of bandwidth settings.

3.5.3 Experiment 3: Trends with varying correlation

From the previous experiment in Section 3.5.2 it was observed that there is a degradation in performance between the uncorrelated and the correlated Rayleigh random sequence. This degradation was observed in the accuracy of the PDF and the accuracy of the higher order moments. The accuracy of the bandwidth was also not sufficient. This experiment investigates the trends in degradation in performance, by varying how strongly correlated the generated samples are. The desired bandwidth is stepped from 0.01 Fs to 0.25 Fs, with all other variables kept constant.

Figure 57 shows the error between the set -3 dB bandwidths and the measured -3 dB bandwidths. Figures 58 and 59 show zoomed in versions of Figure 57 at small bandwidths and relatively large bandwidths. It is clearly visible that at small bandwidths there is about an absolute error of 0.002, which translates to an error of 20% between the smallest -3 dB bandwidth value and the smallest desired -3 dB bandwidth value of 0.01 Fs. As the -3 dB bandwidth becomes larger, the absolute error becomes less. At the larger -3 dB bandwidths the mean of the measured -3 dB bandwidth has an absolute error of less than 0.001. The 99% confidence intervals for the mean at the large -3 dB bandwidths include the desired -3 dB bandwidth values. It is interesting to note that the error for narrow bandwidths is because the measured bandwidth is larger, while the error for larger bandwidths is because the measured bandwidth is smaller, but in practice the absolute error is relatively constant at approximately 0.001 across the entire range.

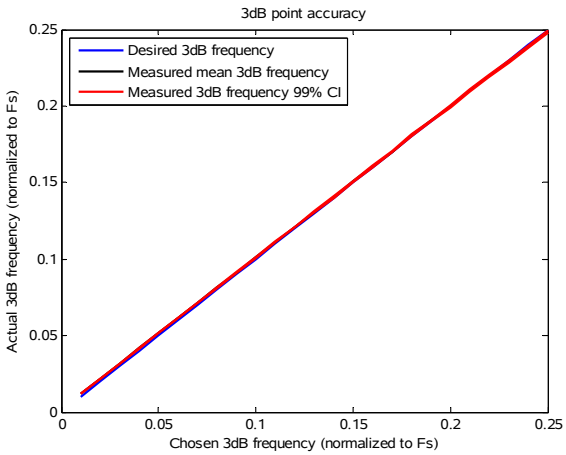

Figure 57: Measured -3 dB bandwidth of the correlated Rayleigh random number.

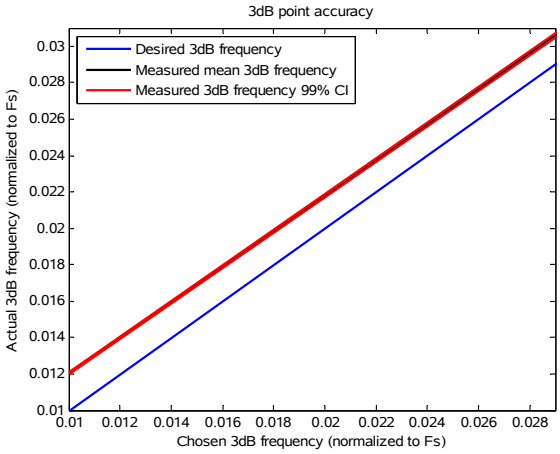

Figure 58: A zoomed in version of the graph of Figure 57 showing small bandwidth settings.

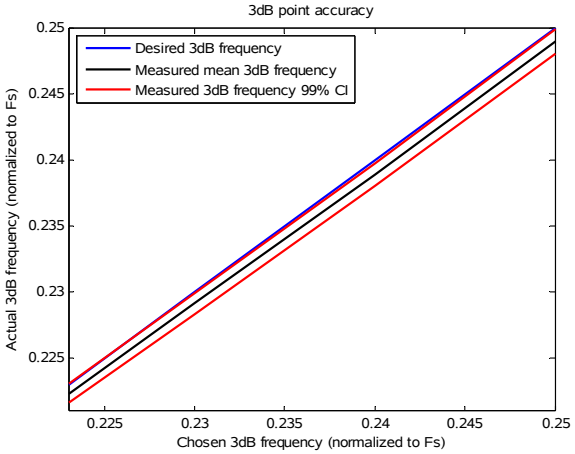

Figure 59: A zoomed in version of the graph of Figure 57 showing large bandwidth settings.

Figure 60 shows the result of χ^2 tests, averaging over 100 experiments for each value of -3 dB bandwidth. Figure 61 shows the number of χ^2 tests out of a 100 that were not rejected for each of the -3 dB bandwidth settings. From Figures 60 and 61 it can be seen that the χ^2 test rejects only a few PDFs from -3 dB bandwidths of 0.05 *Fs* and up. Below this value the χ^2 test rejects an increasing amount of PDFs as the -3 dB bandwidth decreases, until it ultimately rejects more than 55 out of 100. The χ^2 test thus fails to reject the Rayleigh PDF with 99% confidence between bandwidth values of 0.025 *Fs* and 0.25 *Fs* .

Figure 62 shows the KL divergence value with varying degrees of correlation. There is a trend of worsening performance as the correlation in the samples is increased. There appears to be no increase in the variance of the results obtained except for the largest correlation at 0.01 *Fs* . The largest outlier worsens as the correlation is increased as indicated by the red dashed line in the plot. The average result is roughly the same to that of the similar experiment performed on the correlated Gaussian distribution from Section 3.4.4, but as the correlation increases the Rayleigh distribution

performs better than that of the Gaussian distribution. Figure 63 shows the number of samples in the tail of the Rayleigh distribution, summed to give a value out of 100. There is no clear trend, and the number of samples in the tails appear to be constant irrespective of the normalised bandwidth setting. There is however an increase in variability of the result as seen by the variance shown on the plot by the blue dashed lines, and the increase in maximum and minimum outliers shown on the plot by the red lines at the top and the bottom of the graph respectively.

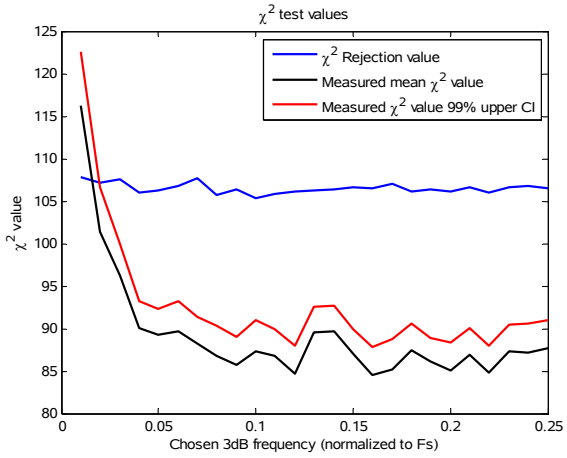

Figure 60: Plot of the outcome from a chi squared test.

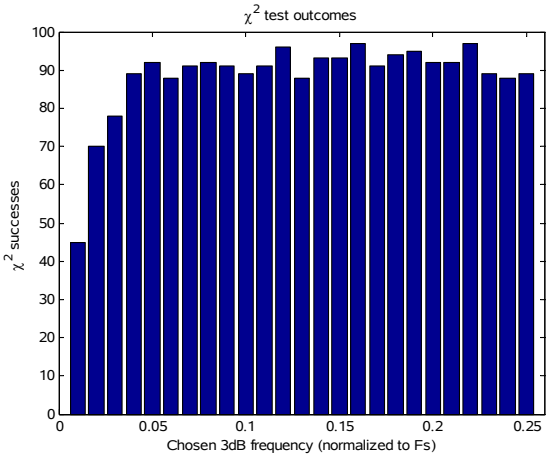

Figure 61: Number of chi squared tests that were unable to reject the PDF.

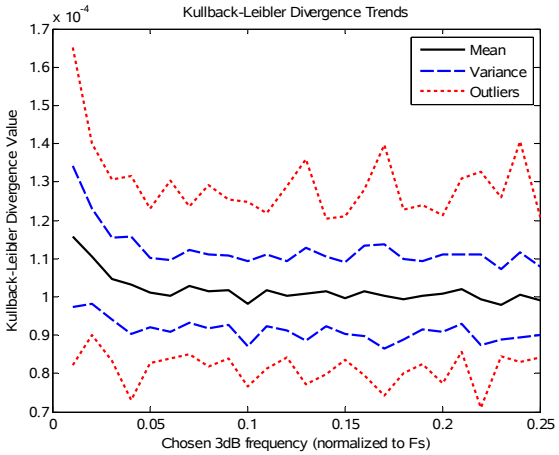

Figure 62: Kullback-Leibler divergence value.

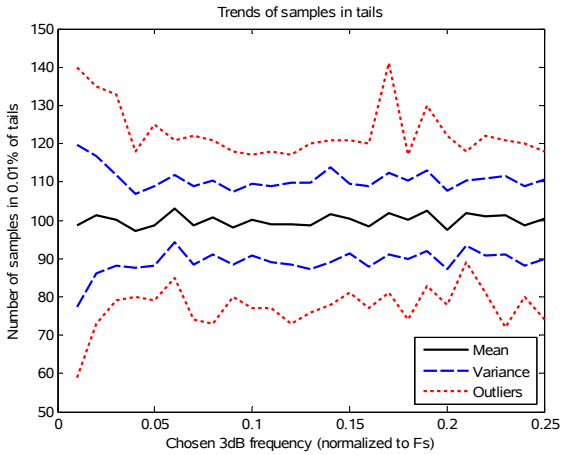

Figure 63: Number of samples in 0.01% of the tail.

Figure 64 shows the results from of the percentage error between the measured higher order moments and the calculated theoretical moments. The trend observed for the χ^2 in Figure 60 is not visible from the higher order moments. All six the moments have the same trend, but there is no clear variation connected with the bandwidth of the generated correlated Rayleigh samples. The larger the order of the moment the larger the percentage error, and the sooner the increase in percentage error due to the smaller bandwidth can be observed. The 6th order moment stayed within the ±0.1% error level for almost all the bandwidth settings between 0.01 Fs and 0.25 Fs.

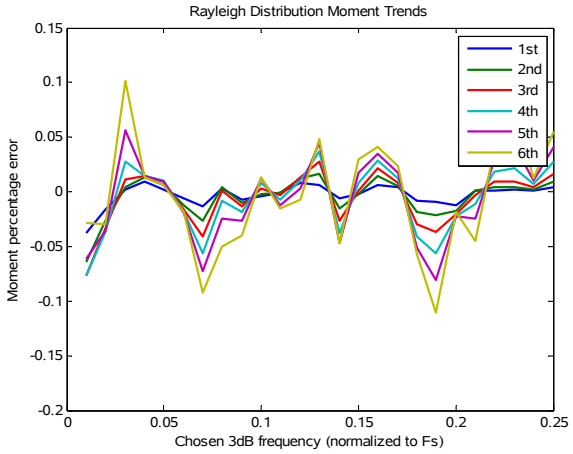

Figure 64: Percentage error of the measured moments from the generated correlated Rayleigh distribution.

This experiment has shown that there is a deviation from the ideal Rayleigh PDF for narrow normalised bandwidths. This deviation stems from the underlying Gaussian random sequences used in the generation of the Rayleigh random sequence. The trend in the even order moments shows a smaller percentage error across the entire tested range of normalised bandwidth settings compared to that of the underlying Gaussian random numbers. The smaller percentage error was also observed in the previous experiment in Section 3.5.2. The α parameter of the Rayleigh distribution influences the scale of the distribution. The performance of the Rayleigh random number sequence is analysed for variation of this parameter.

3.5.4 Experiment 4: Trend with varying Rayleigh scale parameter

This experiment was performed with two independent sequences of Gaussian random numbers as described in the previous two experiments (sections 3.5.2 and 3.5.3). For this experiment the value of the Rayleigh scale parameter α was varied between 0.5 and 3. The bandwidth of the signal was set to 0.025 Fs. One million correlated samples were generated for each of the Gaussian distributed inputs to equation (9), to give one million correlated Rayleigh distributed samples.

The -3 dB accuracy is shown in Figure 65 for values of α between 0.5 and 3. There is a noticeable error of about 6% between the expected value and the measured one. This error was also observed in the previous experiment for the relatively small normalised bandwidth setting of 0.025 Fs. The important result from this figure is that the error remains approximately constant across the entire range of values for the scale parameter α.

Figure 66 shows the result of χ^2 tests, averaged over 100 experiments for each value of α. There is some variation which can be attributed to noise in the χ^2 test due to the variation of the random samples the test was performed on, but there is no trend in the results. The 99% confidence interval remains lower than the threshold for all tested values of α. The χ^2 test thus fails to reject the Rayleigh PDF with 99% confidence for values of α in the tested range between 0.5 and 3.

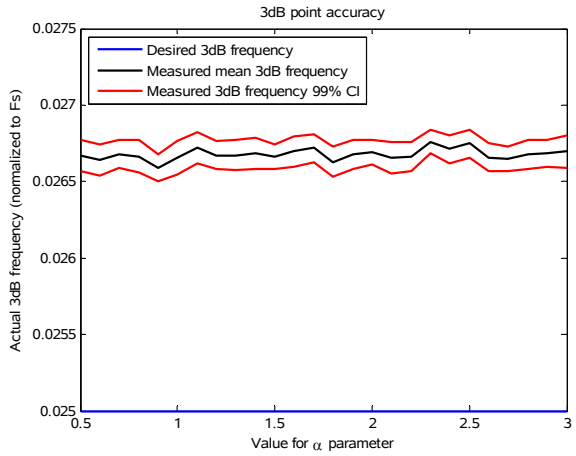

Figure 65: Measured -3 dB bandwidth of the correlated Rayleigh random number.

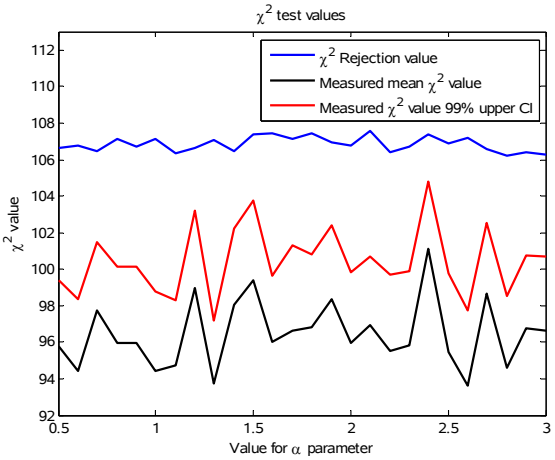

Figure 66: Plot of the outcome from a chi squared test.

Figure 67 shows the number of χ^2 tests out of a 100 that were not rejected for each of the values for α. The average number of tests not rejected is approximately 75. There is no visible trend from this result and no clear correlation between the average χ^2 test values and the number of tests not rejected.

Figure 68 shows the KL divergence value with varying values of the Rayleigh distribution scale parameter. There is no trend in the performance as the scale parameter α is varied. The largest outlier, smallest outlier and variance show no observable trend in the KL divergence value.

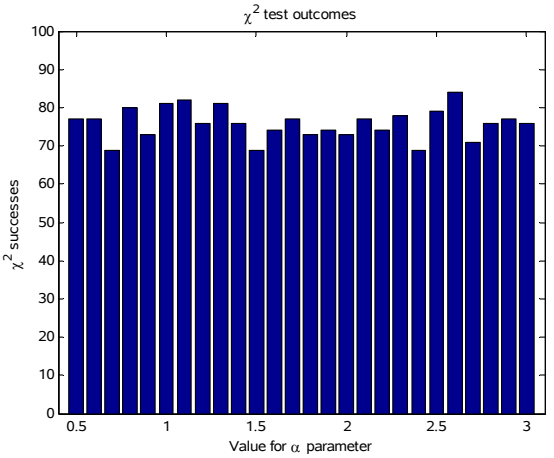

Figure 67: Number of chi squared tests that were unable to reject the PDF.

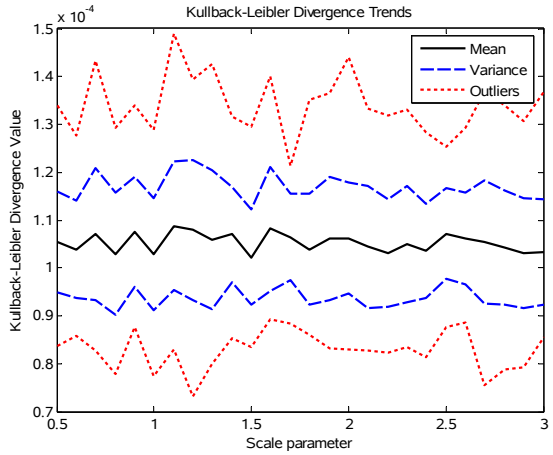

Figure 68: Kullback-Leibler divergence value.

Figure 69 shows the number of samples in the tail of the Rayleigh distribution, summed to give a value out of 100. There is again no trend, and the number of samples in the tails appear to be constant irrespective of the scale parameter setting. Unlike the previous experiment where the correlation was varied, there is no change in variability of the result as seen by the variance shown on the plot by the blue lines.

Figure 70 shows the results from of the percentage error between the measured higher order moments and the calculated ones. Similarly to the χ^2 test the moments reveal no trend in the performance of the generated Rayleigh distribution for all values of α tested.

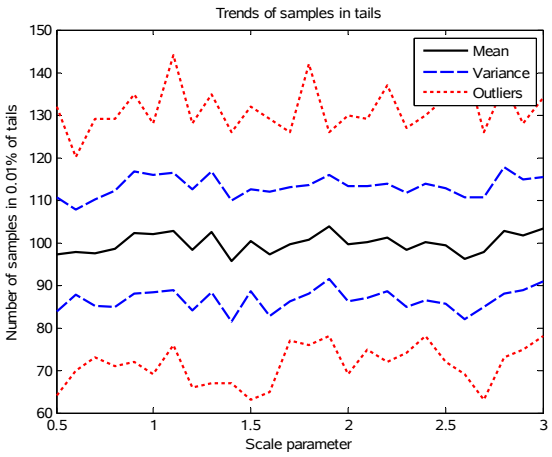

Figure 69: Number of samples in 0.01% of the tail.

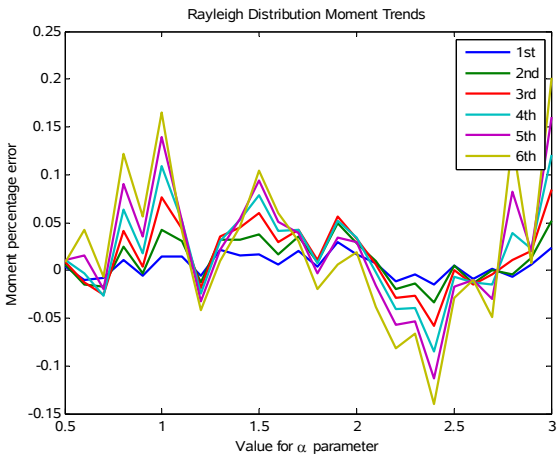

Figure 70: Percentage error of the measured moments from the generated correlated Rayleigh distribution.

The results obtained from this experiment show that there is no measurable variation in the performance of the Rayleigh distributed correlated random sequence when the scale parameter α is varied. This is expected since the scale parameter merely spreads the values over a larger extent, without modifying the shape of the distribution. This is equivalent to multiplying the output of the Rayleigh distributed sequence by a constant, and this linear transformation is not expected to have an influence on performance.

From these results it can be seen that the Rayleigh distributed signal performs well and it is relatively simple to generate. The generation of Log-Normal and Weibull random sequences are more complex and are described in the following sections.

3.6 Log-Normal clutter

The generation of Log-Normal clutter is more complicated than that of Rayleigh distributed clutter. The Log-Normal distribution, equation (14) from Section 2.3.4 repeated here as equation (45), is a distribution that has been found to fit some recorded clutter data measured with high resolution radars.

$$P(x;\sigma,\mu)=\frac{1}{x\sigma\sqrt{2\pi}}e^{-\frac{(\ln(x)-\mu)^2}{2\sigma^2}},x\geq 0 \tag{45}$$

A high range resolution radar is a radar with a high enough range resolution so that the assumption that there are many scatterers in a range bin no longer holds. To obtain a Log-Normal distributed signal a MNLT has to be performed. The parameter μ is an offset value which shifts the distribution to the left or right. A shift in the distribution has been ignored in the experiments that follow by setting $\mu=0$.

First an uncorrelated Log-Normal distributed random sequence will be generated, analysed and compared with previous results. A correlated Log-Normal distribution random sequence will then be generated from correlated Gaussian variables, by applying an MNLT. The PDF error was determined using the χ^2 test and the KL divergence value. Other measures such as the frequency spectrum, phase distribution, autocorrelation function and higher order moments of the resulting random sequences have been investigated.

The method described by [38] was used to generate the Log-Normal random variables. This method is more complex than [39] but it is more accurate. The method described in [39] does not produce a uniform phase distribution (highlighted by [40]), which is required for clutter simulation. These methods map relatively easily to a hardware implementation, which is not the case for spherically invariant random vector (SIRV) type methods. The main drawbacks of the SIRV method lie in the fact that they produce samples in batches, and are computationally more complex than that of MNLT methods [41]. Even though the hardware mapping is relatively simple the calculations are cumbersome and will be relatively time consuming because they would take many clock cycles to perform. This is due to the many multiplications and, with respect to hardware implementation, computationally complex functions (square root, exponential) that have to be implemented. Figure 71 shows the block diagram of the method described in [38].

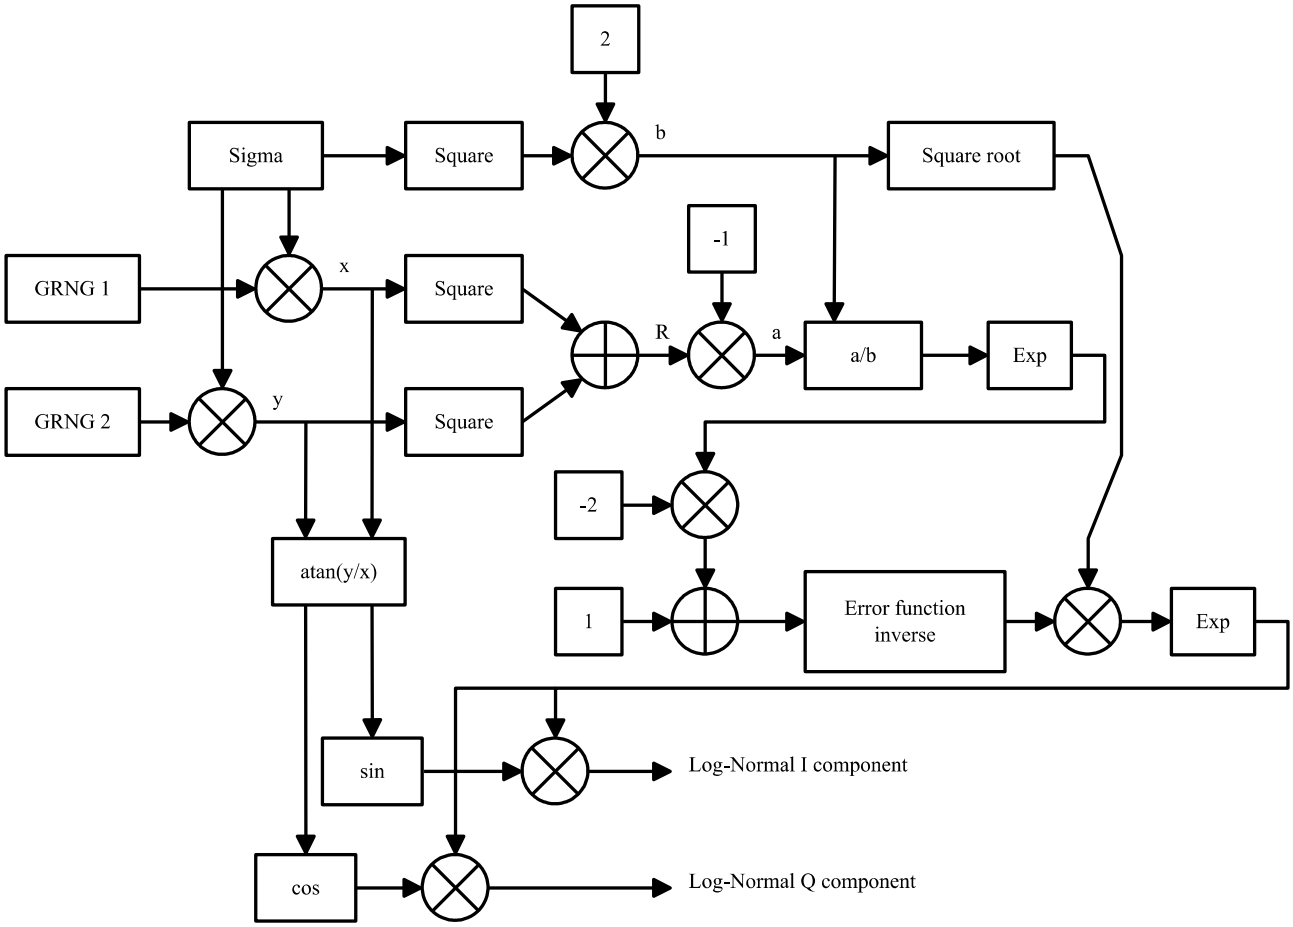

Figure 71: Flow diagram of method proposed in [38] for the generation of a correlated Log-Normal distributed random sequence.

In mathematical terms the flow diagram of Figure 71 can be translated to the following sequence of calculations:

1. Generate 2 independent normalised Gaussian random variables, multiply them with the σ parameter.
2. Calculate R which is given by $R = x^2 + y^2$.
3. Calculate the phase with the relationship $\phi = \tan^{-1}(y/x)$.
4. Use R and σ to calculate the result of $A = \exp\left\{\sqrt{2\sigma^2}\, erf^{-1}[1 - 2\exp((-R)/2\sigma^2)]\right\}$.
5. Finally calculate $Acos(\phi)$ and $Asin(\phi)$, these are the I and Q components of the Log-Normal clutter.

3.6.1 Experiment 1: Generation of the Log-Normal distribution using uncorrelated Gaussian random variables

This experiment serves as the baseline against which the experiments that follow related to the Log-Normal distribution will be compared. To accomplish this a Log-Normal random number sequence has been generated with a trusted method. The value of 0.5 was used for σ. The offset parameter

μ was kept at 0 for all the experiments that follow. There were one million uncorrelated samples generated for each Gaussian distribution, which resulted in one million uncorrelated Log-Normal distributed samples after applying the algorithm from [38].

Figure 72 shows the analytical PDF of the Log-Normal distribution plotted against the histogram of the distribution. There is a visual match between the histogram of the generated samples and the analytical PDF. The mean of the Rayleigh distribution was 1.253 which is only slightly larger than the mean of 1.13 for the Log-Normal distribution, but the last sample of the Rayleigh random sequence was at a value of close to 5.5, while for the Log-Normal distributed random sequence the last sample was at a value close to 12. The reason for this is the much longer tails of the Log-Normal distribution. Because of the longer tails of the Log-Normal distribution the effect of sparse sampling of values in the tail of the distribution is increased. Figures 73 and 74 show the values of the histogram compared to analytical PDF for the leftmost part and the tail respectively. It is fairly straightforward from Figure 74 to see how the sparsity of samples in the tail area could cause very large errors.

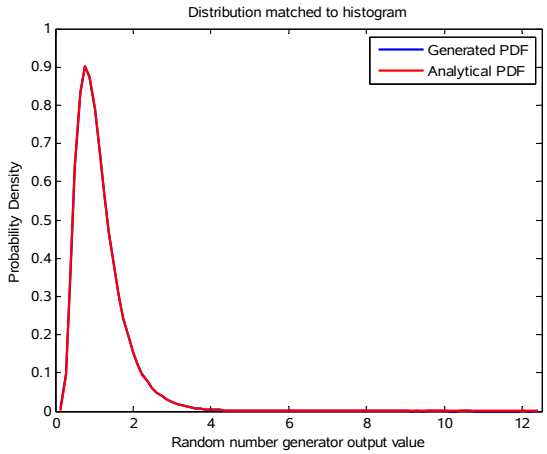

Figure 72: Normalised histogram of the uncorrelated Log-Normal random variable.

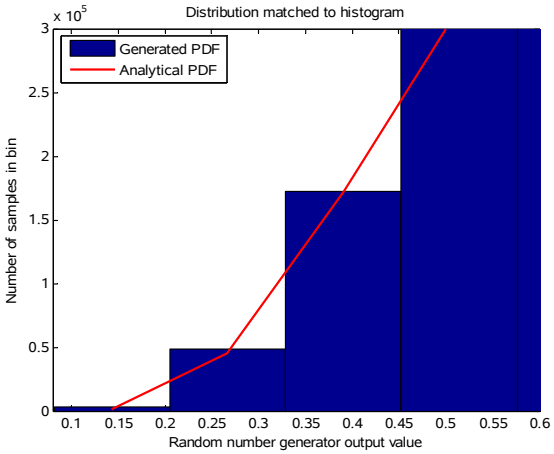

Figure 73: A zoomed version on the start of the generated histogram.

Performing the KL divergence test on the uncorrelated Log-Normal random samples averaged over 30 experiments resulted in a value of 0.000090033. This value is 11.27% less than that of the uncorrelated Gaussian distribution generated in Section 3.4.1.

The moments of order 1 to 6 were calculated and compared to the same order moments measured from the generated data. Figure 75 shows the higher order moments of the analytical Log-Normal distribution plotted against the measured higher order moments from the data. The comparison was done by calculating the percentage error between the measured and calculated moments. The higher order moments are tabulated in Table 13.

Table 13: Percentage error of the higher order moments for the uncorrelated Log-Normal distribution.

Moment order	1st	2nd	3rd	4th	5th	6th
Percentage error	-0.0153%	-0.0599%	-0.1962%	-0.5987%	-1.7240%	-4.5445%

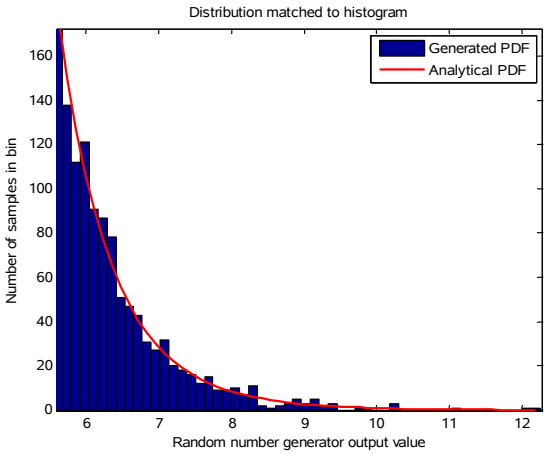

Figure 74: A zoomed version on the tail of the generated histogram.

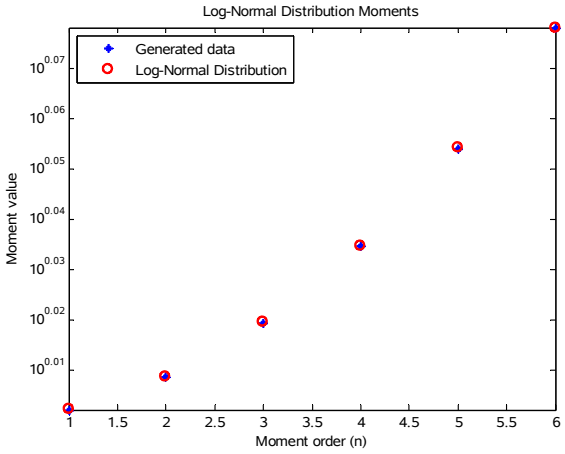

Figure 75: Measured moments from the generated uncorrelated Log-Normal distribution.

The percentage error for the generated Log-Normal moments are much larger than that of the uncorrelated Rayleigh distribution or the uncorrelated Gaussian distributions. Increasing the number of samples reduces the percentage error. This indicates that the method of distribution verification by calculation of the higher order moments is also sensitive to the number of samples used for testing. More samples means that the tails of the distribution are better represented.

Figure 76 shows the PDF of the I channel, and Figure 77 shows the PDF of the Q channel of the Log-Normal random sequence. Both the I and the Q channels have positive and negative values, and they are symmetrical about 0. The I and Q channels also have similar shapes. For the phase of the signal to be uniform from $-\pi$ to π both the I and Q channels have to contain both positive and negative values.

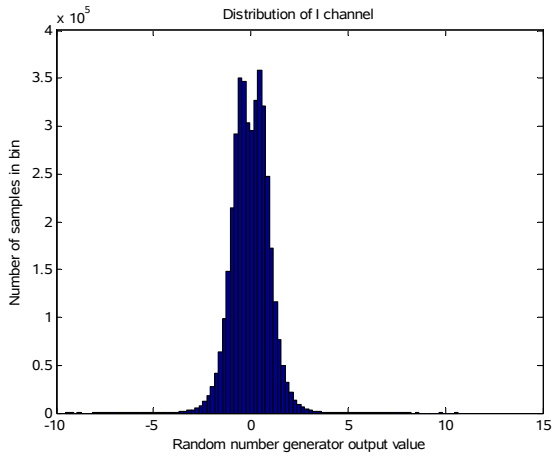

Figure 76: Histogram of the I channel of the algorithm.

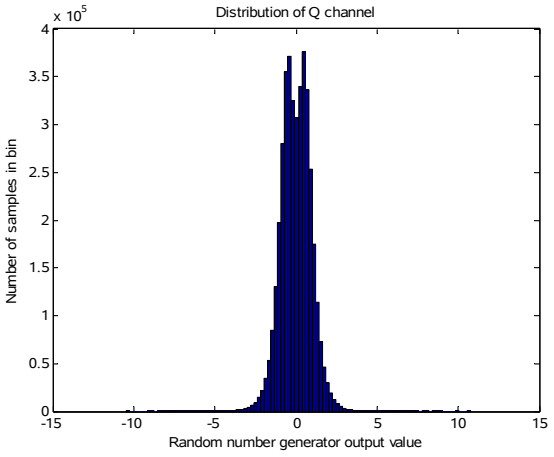

Figure 77: Histogram of the Q channel of the algorithm.

Figure 78 shows the uniform phase distribution of the uncorrelated Log-Normal random sequence. The uniform distribution goes from $-\pi$ to π, which is the correct phase distribution for a clutter signal. Figure 79 shows the frequency spectrum of the Log-Normal distributed random sequence. The frequency spectrum is normalised to sampling frequency with the magnitudes taken relative to the largest frequency component. The FFT is generated from the complex signal, not from the absolute value of the random number sequence. A white spectrum is expected for uncorrelated samples, and this is what is observed.

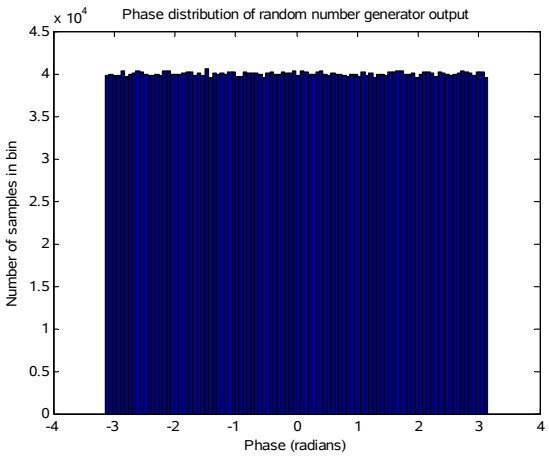

Figure 78: The phase of the uncorrelated Log-Normal random variable.

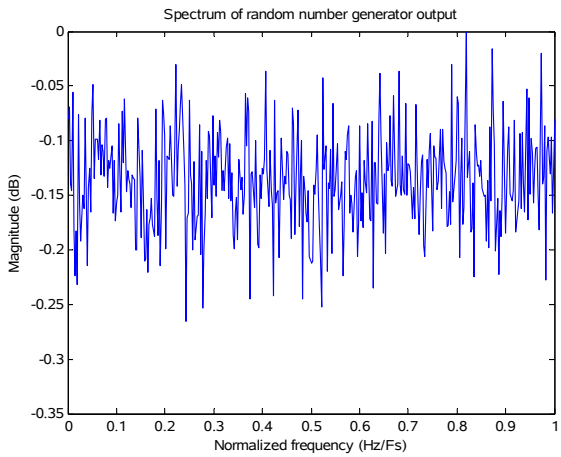

Figure 79: Frequency spectrum of the uncorrelated Log-Normal random variable.

The autocorrelation of the uncorrelated Log-Normal random sequence is shown in Figure 80. There is a spike at the 0 lag sample of the auto correlation, one sample wide, and the spike goes all the way down to approximately 0. This confirms the fact that the consecutive samples of the Log-Normal distribution generated by two independent uncorrelated Gaussian random numbers is also uncorrelated.

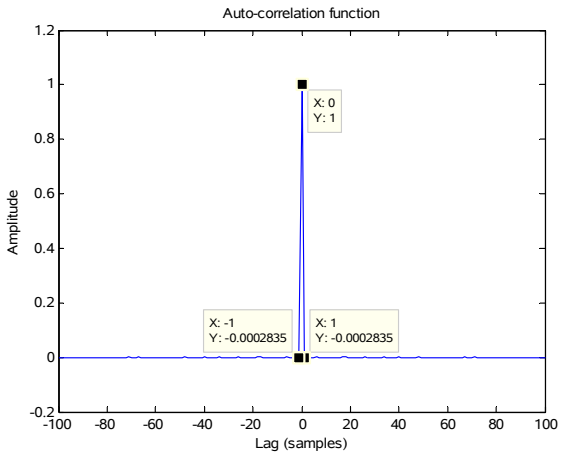

Figure 80: The autocorrelation function of the random variable.

Performing 30 simulations and recording the results of the χ^2, 12 out of the 30 tests were rejected. The mean χ^2 test limit was 139.0, with a 99% lower confidence bound of 131.9. The result from the text on the PDF was 136.2 with 99% confidence intervals at 145.0 and 127.4. Thus on average the χ^2 test could not reject the samples to come from a Log-Normal distribution, but not with 99% confidence as in the previous cases (Rayleigh and Gaussian).

The baseline Log-Normal distribution analysed in this experiment did not perform as well as the baseline Gaussian and Rayleigh distributions. The phase distribution was acceptable for clutter simulation, and there was no underlying correlation in the generated samples, both good properties. The percentage error in the higher order moments were large, but this can be attributed to sparse sampling of the tails, and can be fixed when more amples are generated. The next step is to add a correlation to the random number sequence to analyse its effects on the performance of the distribution.

3.6.2 Experiment 2: Generation of the Log-Normal distribution using two independent time correlated Gaussian random variables

This experiment was performed with two independent sequences of correlated Gaussian random numbers as described with equation (9). Each of the sequences were time correlated with the method discussed in Section 3.2. Each of the Gaussian sequences were passed through the transformation described by the flow diagram in Figure 71. For the experiment one million samples were generated from each of the correlated Gaussian random number generators. Although there was correlation in the sequences within themselves there was no correlation between the two random number sequences, in other words the sequences were independent, just as for the Rayleigh case. After the non-linear transformation has been applied this translated to one million real samples and one million imaginary samples, which together from a complex number sequence of one million samples. For this experiment Fs was set to 0.01 and σ was set to 0.5.

As seen in Figure 81 the resultant measured samples from the generated distribution matches closely to that of the analytical Log-Normal distribution by means of visual inspection. Figure 82 shows the measured bandwidth of the correlated Log-Normal sequence. The frequency spectrum is normalised to sampling frequency with the magnitudes taken relative to the largest frequency component. The value measured at -3 dB was equal to 0.01236, an error of 0.00236, or 23.6%. This error is larger than the error observed for the corresponding experiment performed on the Rayleigh distribution, but smaller than the error observed for the Gaussian.

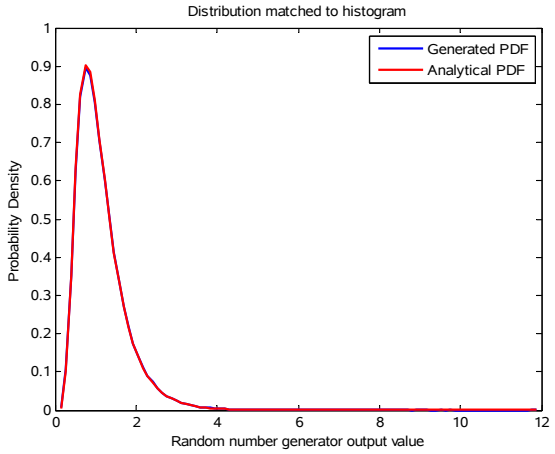

Figure 81: Normalised histogram of the time correlated Log-Normal random variable.

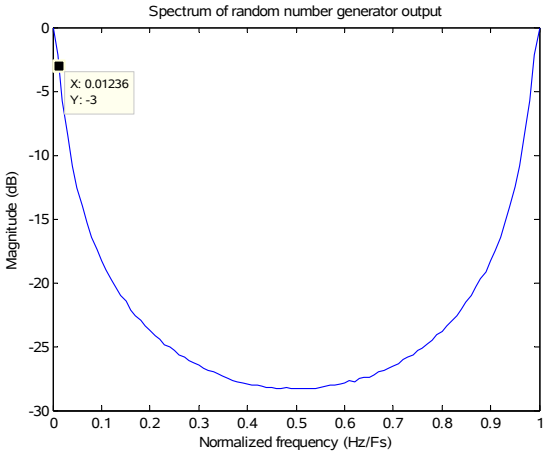

Figure 82: Frequency spectrum of the correlated Log-Normal random variable.

Performing the KL divergence test on the uncorrelated Log-Normal random samples averaged over 30 experiments resulted in a value of 0.0001105. This value is 17.5% less than that of the correlated Gaussian distribution generated in Section 3.4.3. The KL divergence value for the uncorrelated Log-Normal distribution is less than that of the correlated distribution.

The autocorrelation function in Figure 83 shows that there is a correlation in the sequence. The correlation times were measured as 10.9 samples for $\tau_{1/2}$ and 15.8 samples for $\tau_{1/e}$. The signal was a maximum at 0, and completely decorrelated after approximately 100 samples. These results are very similar to the results of the underlying correlated Gaussian sequence shown in Figure 33, discussed in Section 3.4.2. Figure 84 shows that the phase of the correlated Log-Normal sequence is uniform between $-\pi$ and π.

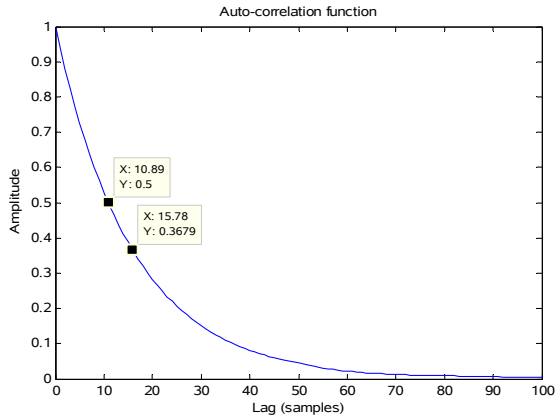

Figure 83: The autocorrelation function of the random variable.

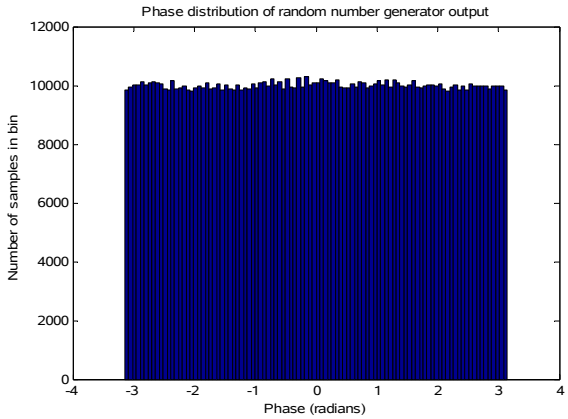

Figure 84: The phase of the correlated Log-Normal random variable.

The higher order moments were recorded for the correlated Log-Normal experiment. The higher order moments are tabulated in Table 14. The percentage error for the correlated Log-Normal distribution is larger than that of the uncorrelated case. The higher the order the larger the

percentage error became, which hints that there are issues in the tails of the distribution. Figure 85 shows the recorded moments plotted against the analytical ones. The figure shows that the measured 6th order moment is markedly less than the analytically calculated 6th order moment, which confirms that the samples in the tail are more sparse than they should be. This is however a single measurement. Because of the statistical nature of the experiment the results will vary.

Table 14: Percentage error of the higher order moments for the correlated Log-Normal distribution.

Moment order	1st	2nd	3rd	4th	5th	6th
Percentage error	0.0301%	-0.0300%	-0.3408%	-1.1760%	-3.1059%	-7.2183%

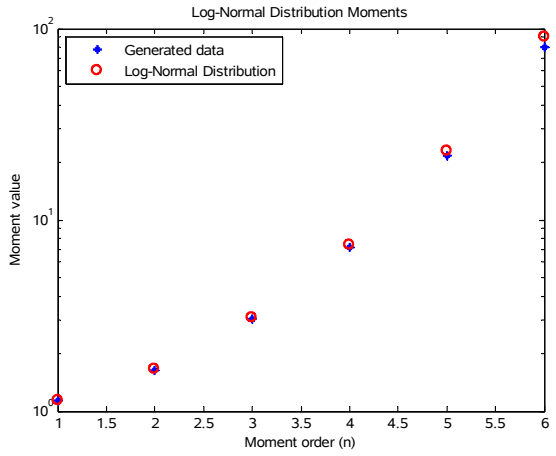

Figure 85: Measured moments from the generated correlated Log-Normal distribution.

For the bandwidth setting of 0.01 Fs, 3 out of 30 χ^2 tests were not rejected. The mean threshold was 141.0, and mean test outcome was 179.0. Thus the generated Log-Normal sequence did not match the analytical Log-Normal distribution with a large enough confidence. Both the percentage error of the higher order moments, and the χ^2 test fails to confirm a satisfactory match between the generated data and a Log-Normal distribution.

The performance of the correlated Log-Normal distribution was only recorded at a single bandwidth setting for this experiment. The next experiment will investigate the performance across a wide range of bandwidth settings.

3.6.3 Experiment 3: Trends with varying correlation

From the previous experiment it was observed that there is a degradation in performance between the uncorrelated and the correlated Log-Normal random sequence. This degradation was observed in the accuracy of the PDF and the accuracy of the higher order moments. This experiment

investigates the trends in degradation in performance, by varying how strongly correlated the generated samples are. The desired bandwidth is stepped from 0.01 Fs in steps of 0.01 to 0.25 Fs, with all other variables kept constant.

Figure 86 shows the measured -3 dB bandwidth and the desired -3 dB bandwidth. Figure 87 shows a zoomed in version of Figure 86 at small bandwidths. It is clearly visible that at small bandwidths there is about an absolute error of 0.002, which translate to an error of 20% between the narrowest -3 dB bandwidth value and the narrowest desired -3 dB bandwidth value. As the -3 dB bandwidth becomes larger, the absolute error becomes less and less. At the larger -3 dB bandwidths the mean of the measured -3 dB bandwidth has an absolute error of less than 0.001. The 99% confidence intervals for the mean at the large -3 dB includes the desired -3 dB bandwidth values.

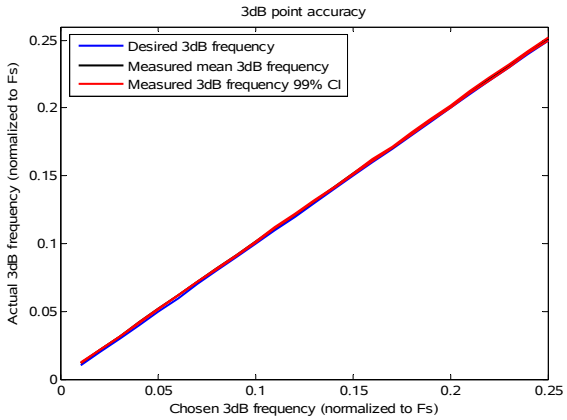

Figure 86: Measured -3 dB bandwidth of the correlated Log-Normal random number.

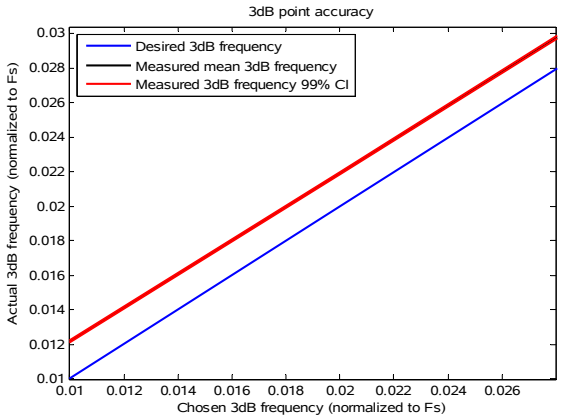

Figure 87: A zoomed in version of the graph of Figure 86 showing small bandwidth settings.

Figure 88 shows the results from the χ^2 tests performed for varying -3 dB bandwidth. The resultant curve shows a similar trend to that of previous tests for the Rayleigh and Gaussian distributions, where the performance of the shape of the PDF degrades the higher the correlation in the samples. The overall performance of the Log-Normal distribution is worse than that of the Rayleigh distribution, because the outcome of the χ^2 test and the 99% confidence interval is in a similar range as the threshold, instead of below it. The poor performance of the generated Log-Normal distribution is also evident from Figure 89. Roughly half of the χ^2 tests performed between 0.03 Fs and 0.25 Fs failed to reject the generated Log-Normal PDF.

Figure 90 shows the KL divergence value with varying degrees of correlation. There is a trend of worsening performance as the correlation in the samples are increased. There appear to be no increase in the variance of the results obtained for a stronger correlation. It is interesting to note that the KL divergence values for the Log-Normal distribution is smaller than that of both the Gaussian and the Rayleigh distributions. This is an interesting results given than the χ^2 test showed poorer

performance. Figure 91 and Figure 92 show the number of samples in the tail of the Log-Normal distribution. For Figure 91 0.01% of the tail is summed to give a value out of 100, and for Figure 92 0.001% of the tail is summed to give a value out of 10. There is no clear trend, and the number of samples in the tails appear to be constant irrespective of the normalised bandwidth setting. There is however an increase in variability of the result as seen by the variance shown on the plot by the blue lines, and the increase in maximum and minimum outliers shown on the plot by the red lines respectively. So even though the mean number of samples in the tail is accurate, the result from a single experiment is less accurate as as the correlation is increased.

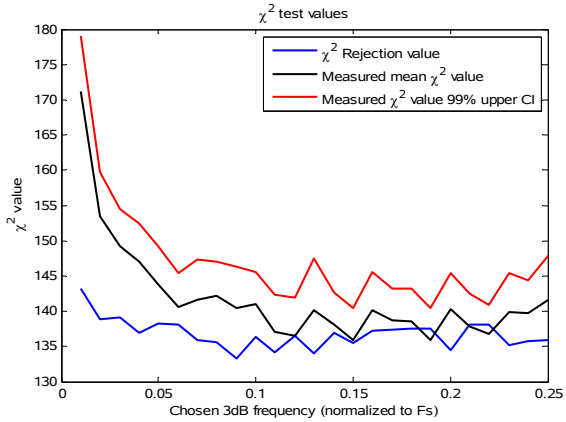

Figure 88: Plot of the outcome from a chi squared test.

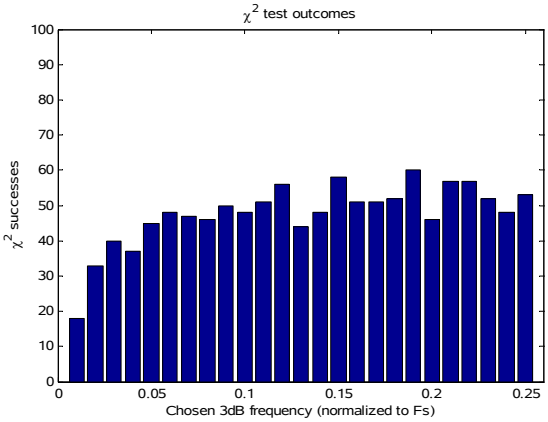

Figure 89: Number chi squared tests that were unable to reject the PDF.

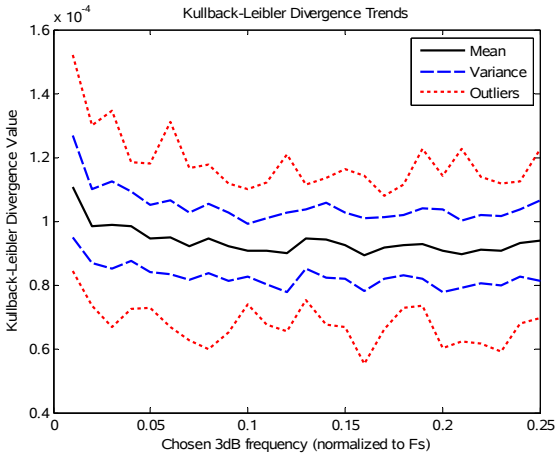

Figure 90: Kullback-Leibler divergence value.

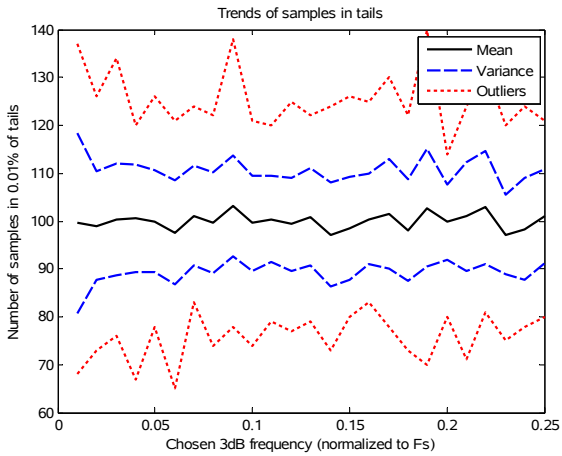

Figure 91: Number of samples in 0.01% of the tail.

Figure 93 shows the results from of the percentage error between the measured higher order moments and the calculated ones. All size the measured moments have the same trend with an error varying around zero percent irrespective of the bandwidth setting. The larger the order of the moment the larger the percentage error. The only visible trend is in the varying size of the errors which increases for smaller bandwidth settings, which was also observed with the KL divergence value.

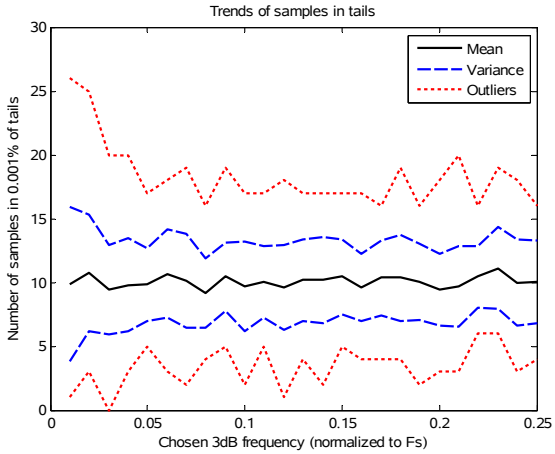

Figure 92: Number of samples in 0.001% of the tail.

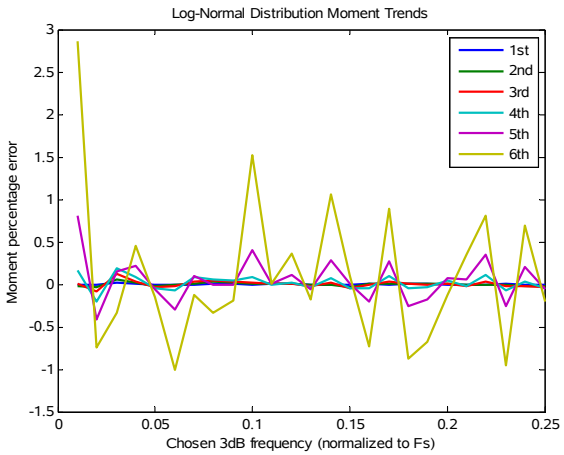

Figure 93: Percentage error of the measured moments from the generated correlated Log-Normal distribution.

This experiment has shown that the generated distribution struggles to be matched to the analytical Log-Normal distribution. There is a deviation from the ideal Log-Normal PDF for all bandwidth settings, and this deviation increases when the normalised bandwidths are set to small values. The trend in deviation from the ideal PDF stems from the underlying Gaussian random sequences used in the generation of the Log-Normal random sequence. All the tests except the KL divergence test showed poorer performance for the Log-Normal distribution than the others previously examined. The reason for the overall poor performance is unknown, and presumably stems from the usual problem of sparse sampling in the tails, the effect of which is exaggerated due to the relatively long tails of the Log-Normal PDF.

3.6.4 Experiment 4: Trend with varying Log-Normal scale parameter

This experiment was performed with two independent sequences of Gaussian random numbers as described in the previous two experiments. For this experiment the value of the Log-Normal scale parameter σ was varied between 0.05 and 1 in steps of 0.05. Even though σ is known as the scale parameter for the Log-Normal distribution, it also affects the shape of the distribution. The bandwidth of the signal was set to 0.025 Fs. One million correlated samples were generated for each of the Gaussian distributed inputs to the algorithm of Figure 71, to give one million correlated Log-Normal distributed samples.

Figure 94 shows the trend of the average χ^2 value calculated from the error of the Log-Normal distribution. The scale of Figure 94 does not allow for a thorough inspection of the result but it is clear that between $\sigma=0.5$ and $\sigma=0.6$ the performance of the PDF shape formed from the generated samples degrades drastically. Figure 95 is a zoomed in plot of Figure 94. From this figure it can be observed that the χ^2 test fails to reject the PDF of the generated Log-Normal random

sequence for values of sigma between 0.05 and 0.5, but from there onward all the way up to $\sigma=1$ the χ^2 test rejects the generated PDF. The value of $\sigma=0.5$ is the crossover point. This explains the marginal performance observed for experiment 3 since σ was held constant at 0.5 to measure the performance of the bandwidth.

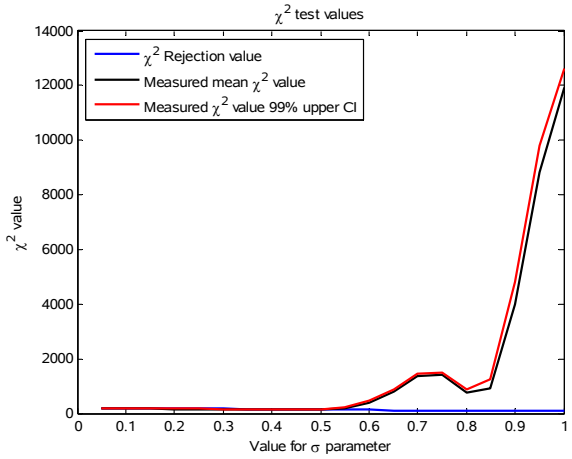

Figure 94: Plot of the outcome from a chi squared test.

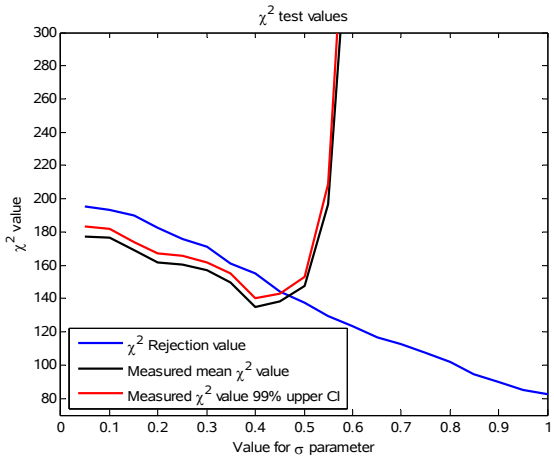

Figure 95: Plot zoomed in version of Figure 94 focusing on the trend of the chi-squared test rejection value.

From Figure 96 it can be seen that after $\sigma=0.5$ the χ^2 tests start to reject the generated samples as a Log-Normal PDF. Just over 50 tests were not rejected between $\sigma=0.05$ and $\sigma=0.45$ as can be seen from the histogram, and on average the χ^2 test value exceeded the χ^2 rejection threshold. The χ^2 rejection threshold level decreases with an increase in σ. This is due to the number of bins in the measured histogram of the PDF. The Log-Normal distribution's tails becomes longer with respect to the body of the distribution as σ increases. This influences the ratio between the number of bins that capture the samples in the body, and the number of bins that capture the samples in the tails. The total number of bins decreases because samples in the tail of the distribution are summed to minimise the effect of their relatively large error on the χ^2 test outcome (as discussed in Section 3.2 on page 30 [30]). This strategy has not worked for this experiment.

Figure 97 shows the KL divergence value with varying values of the Log-Normal distribution scale parameter. There is a trend indicating improved performance for larger values of the scale parameter. The largest outlier, smallest outlier and variance show a slight increase in variability of the data as the scale parameter is increased.

Figure 98 and Figure 99 show the number of samples in the tail of the Log-Normal distribution. For Figure 98 0.01% of the tail is summed to give a value out of 100, and for Figure 99 0.001% of the tail is summed to give a value out of 10. There is no clear trend, and the number of samples in the tails appear to be constant irrespective of the Log-Normal scale parameter setting.

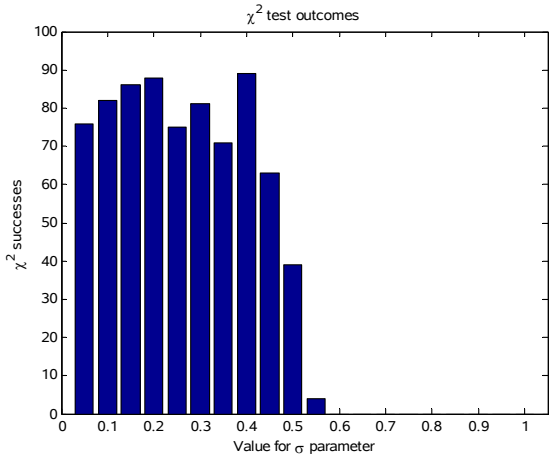

Figure 96: Number of the chi squared tests that were unable to reject the PDF.

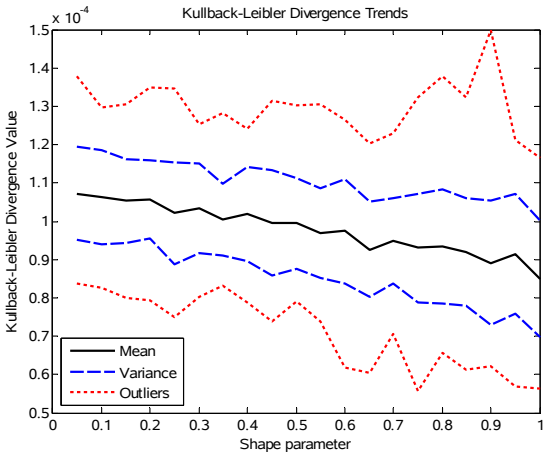

Figure 97: Kullback-Leibler divergence value.

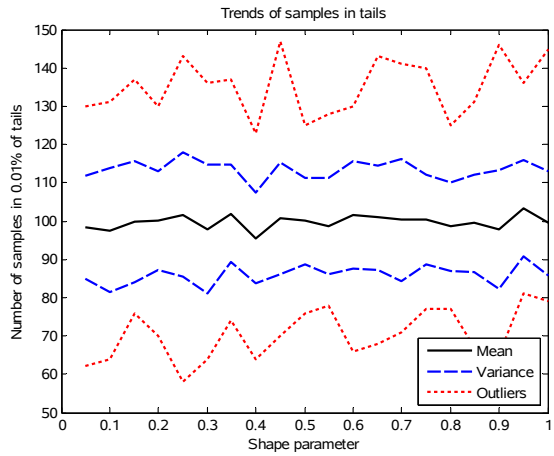

Figure 98: Number of samples in 0.01% of the tail.

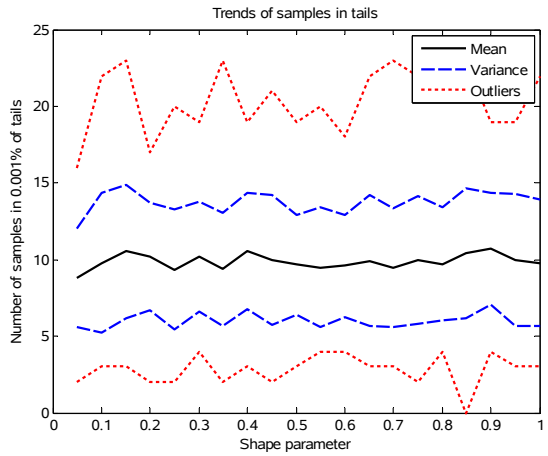

Figure 99: Number of samples in 0.001% of the tail.

The full scale of the percentage error for the higher order moments are shown in Figure 100 and a zoomed in version between -1% and 1% error is shown in Figure 101. The percentage error increases as σ increases. At approximately $\sigma=0.7$ the 6th order moment percentage error declines sharply and continues to do so all the way to $\sigma=1$. The percentage error in the moments does not show the same sharp switch between acceptable and poor performance that the χ^2 test has. The χ^2 test showed a clear change in performance at approximately $\sigma=0.5$, where as the same degradation is observed at $\sigma=0.7$ with the higher order moments. The percentage error of the moments do however indicate that the tail of the generated Log-Normal distribution is the source of the problems regarding the accuracy of the distribution, since the 5th and 6th order moments have shown the largest errors.

Figures 102, 103 and 104 show histograms of the generated samples with the analytical Log-Normal distribution plotted over them for comparison. Figure 102 was generated with $\sigma=0.1$, Figure 103 with $\sigma=0.6$ and Figure 104 with $\sigma=0.9$. The higher the value for σ the larger the tail of the Log-Normal distribution becomes. For $\sigma=0.1$ there are no significant tails on either side

of the distribution, but as σ increases to $\sigma=0.6$ the tail starts to become significant, stretching over a range from 0 to approximately 18. When $\sigma=0.9$ the tail has increased to include the entire range up to 137, which is extreme.

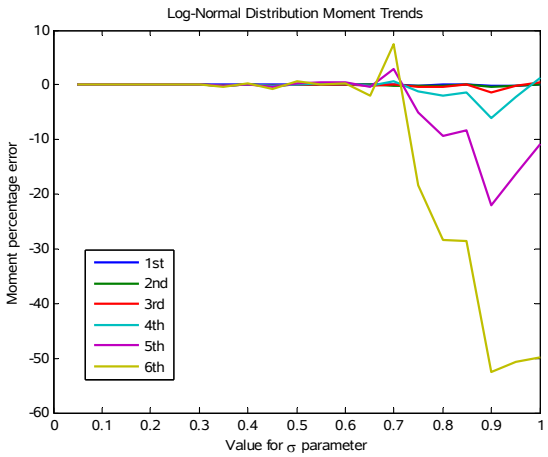

Figure 100: Percentage error of the measured moments from the generated correlated Log-Normal distribution.

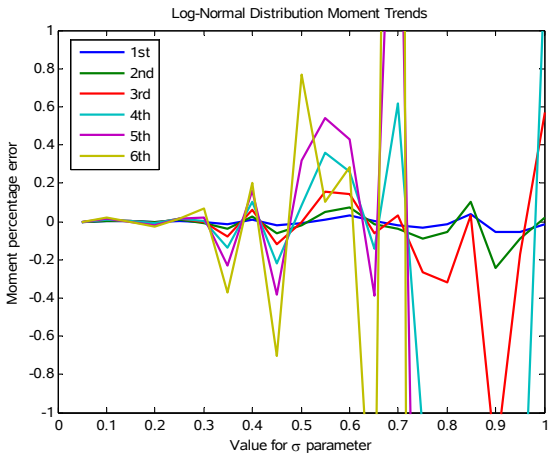

Figure 101: A zoomed version of the percentage error of the measured moments from the generated correlated Log-Normal distribution.

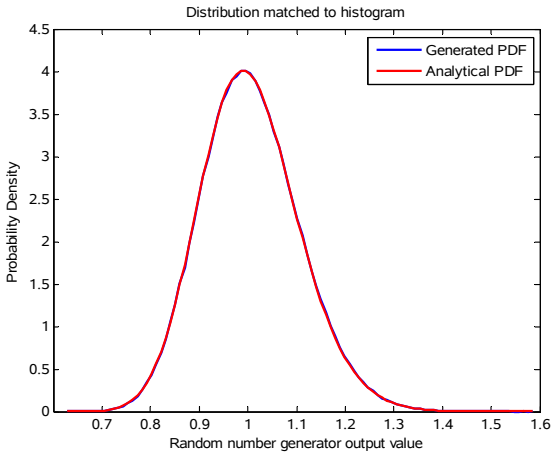

Figure 102: Normalised histogram of the Log-Normal random variable for sigma equal to 0.1.

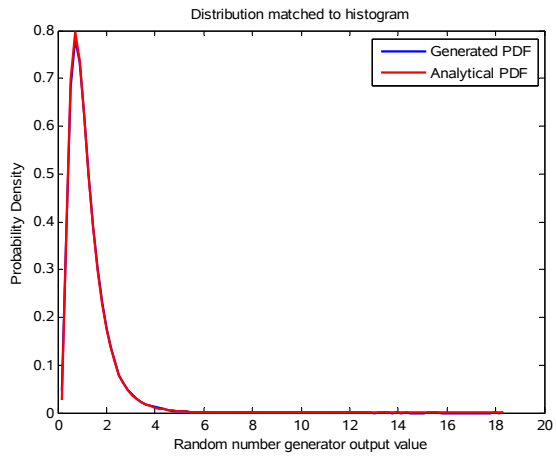

Figure 103: Normalised histogram of the Log-Normal random variable for sigma equal to 0.6.

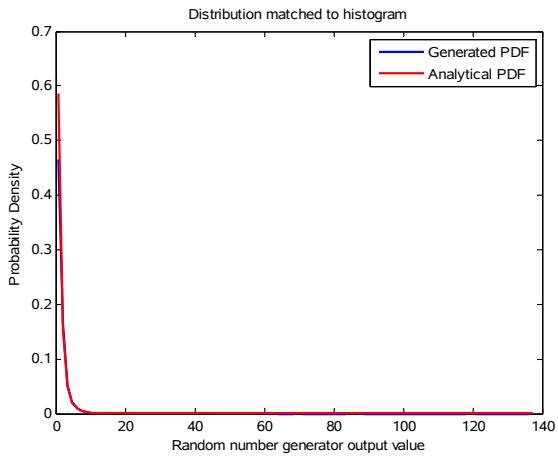

Figure 104: Normalised histogram of the Log-Normal random for sigma equal to 0.9.

One million samples were generated, but only 10 of those fall between 45 and 100 as shown in Figure 105, the other 999990 fall between 0 and 45. Substituting σ equation (14) with 0.9 one obtains equation (46). Integrating from 45 to infinity produces the result in equation (47).

$$P(x;0.9,0)=\frac{1}{0.9x\sqrt{2\pi}}e^{-\frac{(\ln(x))^2}{2(0.9)^2}},x\geq 0 \tag{46}$$

$$\int_{45}^{\infty}\frac{1}{0.9x\sqrt{2\pi}}e^{-\frac{(\ln(x))^2}{2(0.9)^2}}=0.0000117041 \tag{47}$$

The area of the histogram is equal to one million because each sample has an area of 1. The result from equation (47) has to be multiplied by this area of one million so a comparison of the actual area and the required area for that part of the tail can be made. When this multiplication is performed the result is 11.7 samples, which is only 1.7 samples larger than the 10 generated for the tail. The main problem thus does not necessarily lie with the number of samples in the tail, but the positioning and fidelity of those samples. It is highly unlikely for the randomly generated samples to lie at exactly the correct position to produce an accurate moment result. The samples themselves have a fixed area of 1, which has to be placed at a single position, even though ideally that area has to be spread out over a larger area. Increasing the number of generated samples can solve this problem, but increasing the number of samples by a factor of 100 will require the generation of more than a hundred million samples. This will result in the use of approximately 3 Gigabytes of memory just to store the samples, ignoring the memory required to run any form of analysis, and ignoring the memory overhead required for the operating system and the program. This is a somewhat impractical requirement for today's desktop computers. But even if this tasks was performed the problem will only be shifted further into the tails, not resolved.

Figure 106 shows the variation in bandwidth as a function σ. The bandwidth was set to 0.025 *Fs* for all values of σ. The bandwidth measured here shows a trend where the measured bandwidth increases, nearing the set bandwidth of 0.025 *Fs* as σ becomes larger. The narrowest bandwidth measured was 0.0158 *Fs* at $\sigma=0.05$ and the widest bandwidth measured was 0.0223 *Fs* at $\sigma=1$. These values are more than ten times less than the required bandwidth, and it seems that the shape of the Log-Normal distribution has very little effect on the accuracy of the bandwidth with the clutter generation algorithm implemented from [38].

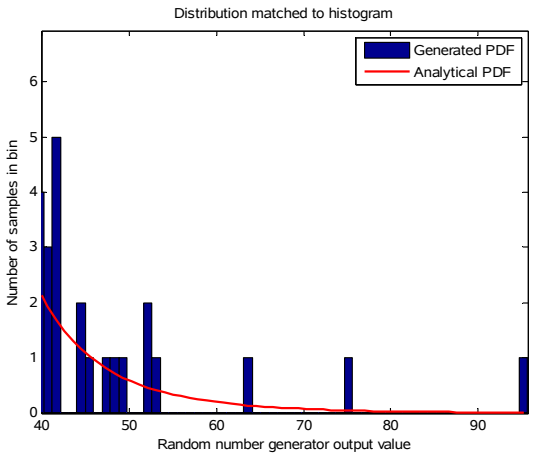

Figure 105: A zoomed version on the tail of the Log-Normal histogram for sigma equal to 0.9.

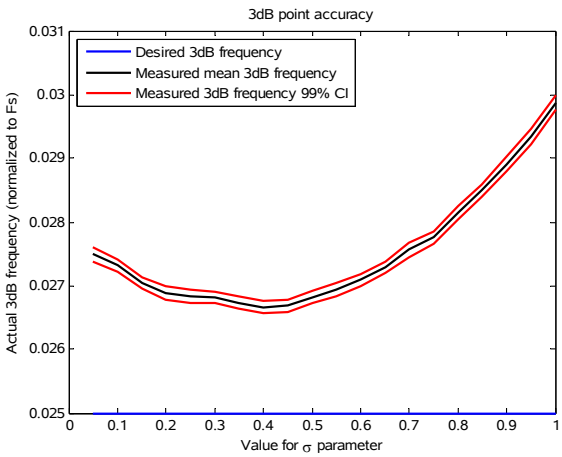

Figure 106: Measured -3 dB bandwidth of the correlated Rayleigh random number.

From these results it can be seen that the Log-Normal distributed random sequence generated by the method described in [38] performs reasonably well, but does suffer from sparse sampling effects in the tails. This is not a problem of the method but rather a problem which highlights the difficulty of accurately evaluating a distribution with long tails. This random sequence generation method is complex and would require a large implementation effort for use on an FPGA based hardware clutter simulation system. The method described in [38] can be modified to produce Weibull distributed clutter, but due to the complexity of the algorithm, an alternative and less complex algorithm from [42] was investigated.

3.7 Weibull clutter

The generation of Weibull clutter is less complicated than the method used for the Log-Normal distribution in the previous section, but more complicated than the method used for the Rayleigh distributed samples. The Weibull distribution, equation (19) from Section 2.3.7 repeated here as equation (48), has been fitted to recorded clutter data measured with high resolution radars.

$$P(x;\lambda,k)=\frac{k}{\lambda}\left(\frac{x}{\lambda}\right)^{k-1}e^{-\left(\frac{x}{\lambda}\right)^{k}},\ x\geq 0 \tag{48}$$

An MNLT is performed on Gaussian distributed samples to obtain Weibull distributed samples. Unlike the Log-Normal and Gaussian distributions the Weibull distribution does not have a parameter to change its position. The parameters of the Weibull distribution are known as the scale and shape parameters, denoted by λ and k respectively.

First an uncorrelated Weibull distributed random sequence was generated, analysed and compared with the uncorrelated results from previously analysed distributions. A correlated Weibull distribution random sequence was then generated from correlated Gaussian variables, by applying

an MNLT. The PDF error, KL divergence value, frequency spectrum, phase distribution, autocorrelation function, χ^2 error and percentage error of the higher order moments of the resulting random sequences was then analysed.

The method described by [42] was used to generate the Weibull random variable. This method is less complex than the one in [38] which was used to generate the Log-Normal variable discussed in the previous section. Figure 107 shows the block diagram of the method described in [42].

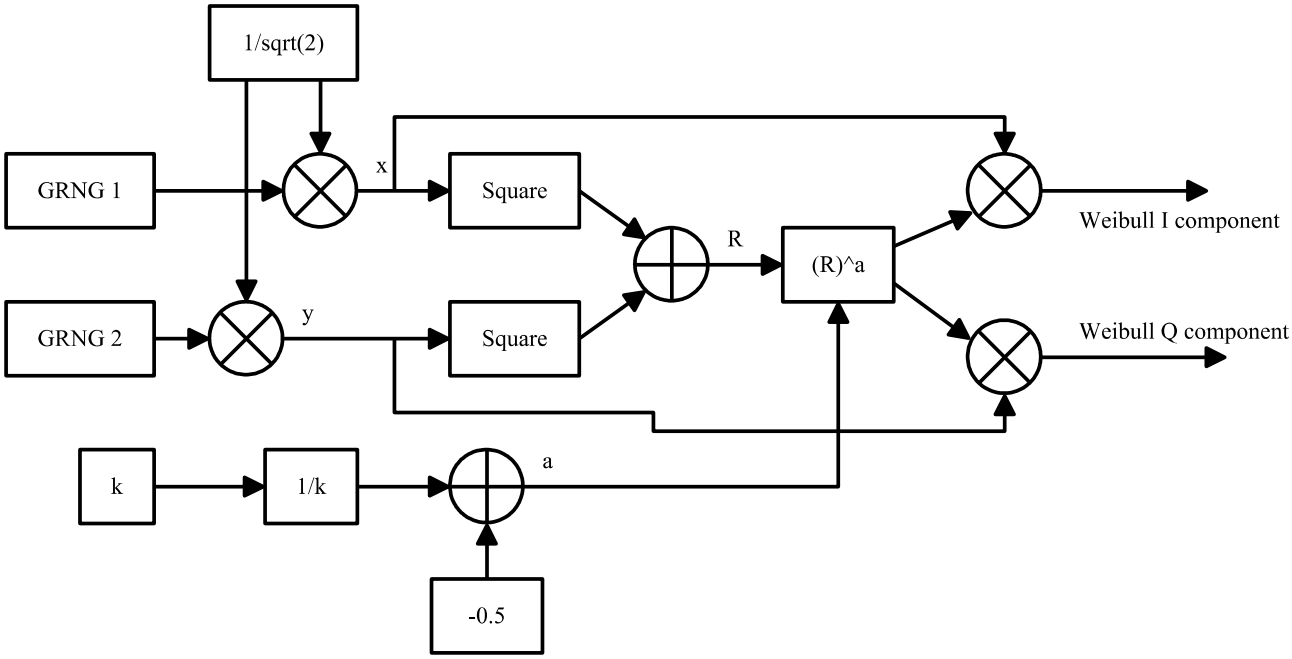

Figure 107: Flow diagram of method proposed in [42] for the generation of a correlated Weibull distributed random sequence.

In mathematical terms the flow diagram of Figure 107 can be translated to the following sequence of calculations:

1. Generate 2 independent normalised Gaussian random variables, and multiply them with $1/\sqrt{2}$.
2. Calculate $R=x^2+y^2$.
3. Calculate the parameter a using $a=(1/k-1/2)$.
4. The I and Q components of the resulting Weibull random sequence are simply xR and yR respectively.

3.7.1 Experiment 1: Generation of the Weibull distribution using uncorrelated Gaussian random variables

This experiment serves as the baseline against which the experiments that follow related to the Weibull distribution will be compared. To accomplish this a well behaved Weibull random number sequence has been generated, with a value of 1.5 for the shape parameter k. The Weibull distribution reduces to the Rayleigh distribution for the shape parameter k set equal to 2. Since it

was observed with the Rayleigh distribution that the scale parameter does not affect the results, the Weibull scale parameter λ was kept at 1 for all the experiments that follow. One million uncorrelated samples generated for each Gaussian distribution, which resulted in one million uncorrelated Weibull distributed samples after applying the algorithm from [42].

Figure 108 shows the analytical PDF of the Weibull distribution plotted against the histogram of the distribution. There is a visual match between the histogram of the generated samples and the analytical PDF. There is no significant deviation from the analytical PDF by the generated PDF up to about $x=3.5$. At this point there are less than 250 samples in a bin and this number decreases drastically to less than 15 samples in a bin by $x=4.5$ as can be seen in Figure 109. These results are similar to the uncorrelated Log-Normal and uncorrelated Rayleigh distribution, with a degradation in performance of the PDF shape in the tails.

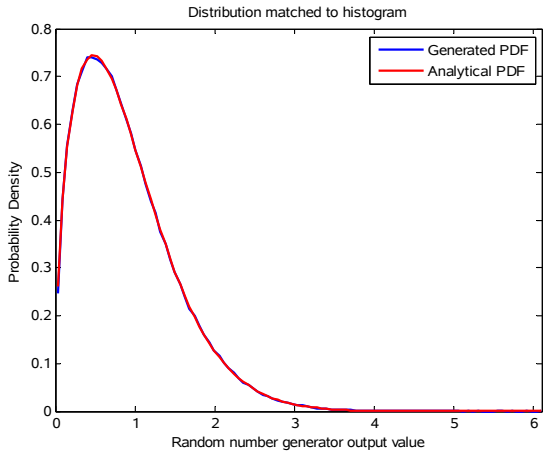

Figure 108: Normalised histogram of the uncorrelated Weibull random variable.

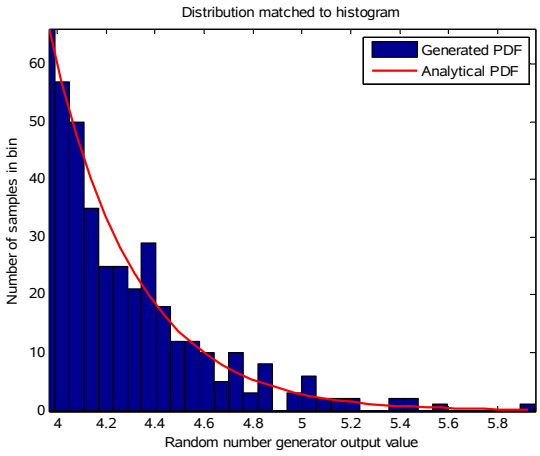

Figure 109: A zoomed version on the tail of the generated histogram.

Performing the KL divergence test on the uncorrelated Weibull random samples averaged over 30 experiments resulted in a value of 0.00009757. This value is 3,84% less than that of the uncorrelated Gaussian distribution generated in Section 3.4.1.

The moments of order 1 to 6 were calculated and compared to the same order moments measured from the generated data. Figure 110 shows the theoretical higher order moments of the analytical Weibull distribution plotted against the measured higher order moments from the data. The comparison was accomplished by calculating the percentage error between the measured and calculated moments. The higher order moments are tabulated in Table 15. This small percentage error is an excellent result since even the error in the 6th order moment (which is the most sensitive moment) was very close to 1%. The percentage error for the generated Weibull moments are much smaller than those of the uncorrelated Log-Normal distribution but larger than the percentage error for the uncorrelated Rayleigh and the uncorrelated Gaussian distributions.

Table 15: Table 15: Percentage error of the higher order moments for the uncorrelated Weibull distribution.

Moment order	1st	2nd	3rd	4th	5th	6th
Percentage error	0.0088%	0.0527%	0.1794%	0.3765%	0.6569%	1.0700%

Figure 111 shows the frequency spectrum of the Weibull distributed random sequence. The spectrum is normalised to sampling frequency with the magnitudes taken relative to the largest frequency component. The FFT is generated from the complex signal, not from the absolute value of the random number sequence. A white spectrum is expected for uncorrelated samples and this is indeed the result observed.

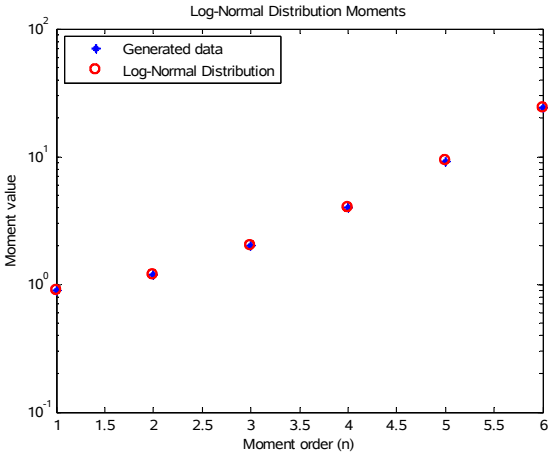

Figure 110: Measured moments from the generated uncorrelated Weibull distribution.

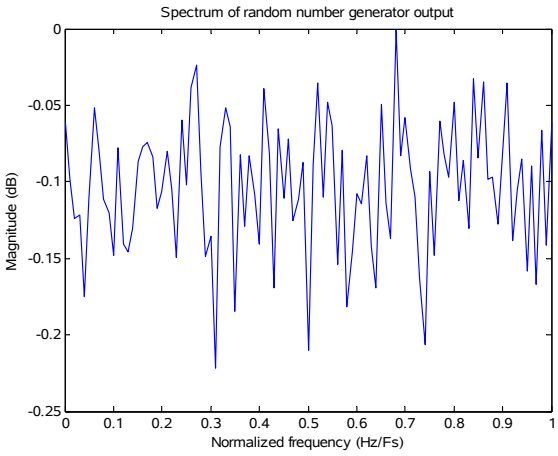

Figure 111: Frequency spectrum of the uncorrelated Weibull random variable.

Figure 112 shows the result from autocorrelation. There is a spike at the 0 lag sample of the auto correlation, one sample wide, and the spike goes all the way down to approximately 0. This confirms the fact that the consecutive samples of the Weibull distribution generated by two independent uncorrelated Gaussian random numbers is also uncorrelated.

The histogram of the phase of the uncorrelated Weibull random number sequence in Figure 113 shows a uniform distribution from $-\pi$ to π. This is the desired result for the phase distribution. A uniformly distributed phase is required for a clutter signal and the method used in this section adheres to this requirement.

Figure 114 shows the PDF of the I channel, and Figure 115 shows the PDF of the Q channel of the Weibull random sequence. Both the I and the Q channels have positive and negative values, and they are symmetrical about 0. The I and Q channels also have similar shapes.

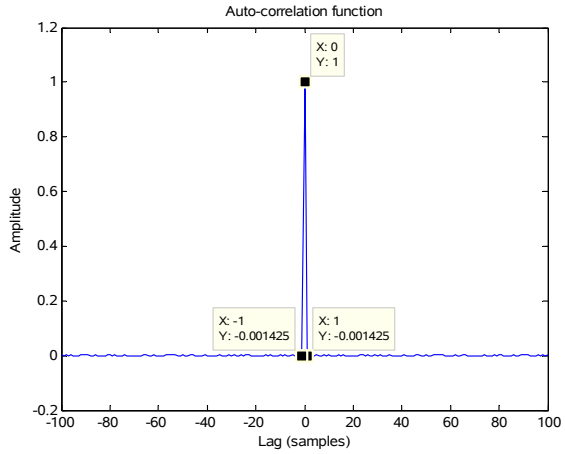

Figure 112: The autocorrelation function of the random variable.

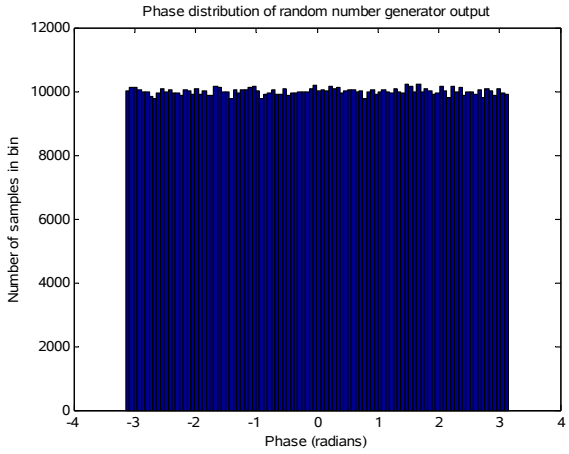

Figure 113: The phase of the uncorrelated Weibull random variable.

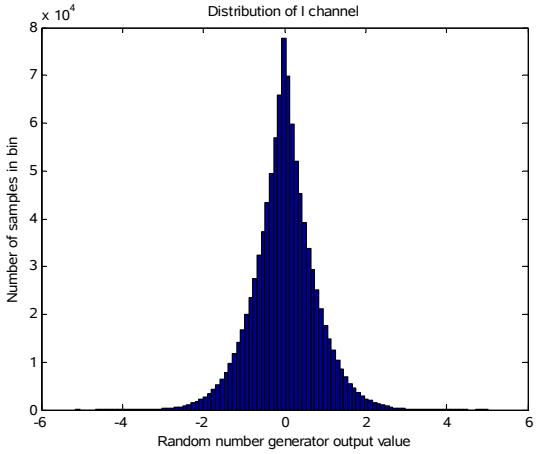

Figure 114: Histogram of the I channel of the algorithm.

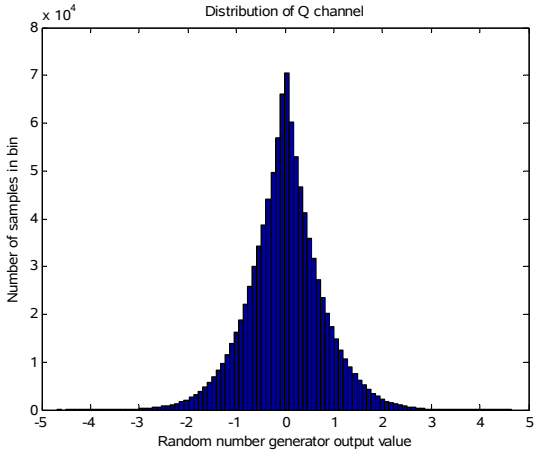

Figure 115: Histogram of the Q channel of the algorithm.

Repeating the experiment 30 times and performing a χ^2 test provides a measure of how well the PDF of the generated Weibull distribution fits to the required Weibull PDF. The χ^2 test was not rejected 22 times out of the 30 and on average the test result was not rejected with a confidence of more than 99%.

The baseline Weibull distribution analysed in this experiment performed well, and performed on the same level as that of the Rayleigh and Gaussian distributions. This hints that the results that follow may be favourable, since the results for the Rayleigh and Gaussian distributions were positive. The next step is to add a correlation to the random number sequence to analyse its effects on the performance of the generation technique.

3.7.2 Experiment 2: Generation of the Weibull distribution using two independent time correlated Gaussian random variables

This experiment was performed with two independent sequences of correlated Gaussian random numbers as described with equation (9). Each of the sequences were time correlated with the method discussed in Section 3.2. Each of the Gaussian sequences was transformed with the process described by the flow diagram in Figure 107 from the start of Section 3.7. For this experiment one million samples were generated from each of the correlated Gaussian random number generators. Although there was correlation in the sequences with themselves there was no correlation between the two random number sequences, in other words the I and Q sequences were independent. After the non-linear transformation has been applied from [42] this translated the one million real and one million imaginary samples, into a complex number sequence of one million samples. For this experiment Fs was set to 0.01 and k was set to 1.5.

Figure 116 shows the analytical PDF of the Weibull distribution plotted over the correlated random samples generated as described in the previous paragraph. The analytical Weibull PDF matches closely to that of the generated samples by means of visual inspection. The KL divergence value obtained when averaging over 30 experiments was 0.00011313. This value is 11.49% greater than that of the uncorrelated case obtained for the previous experiment. Because a smaller value is better, this is an indication that when correlation is added to the random sequence the performance could be degraded. A more through investigation of this result is required before a valid conclusion can be made.

Figure 117 shows the measured bandwidth of the correlated Weibull sequence. The frequency spectrum is normalised to sampling frequency with the magnitudes taken relative to the largest frequency component. The value measured at -3 dB was equal to 0.01217, an error of 0.00217, or 21.7%. This result is similar to that observed for the previous methods and even the underlying Gaussian distribution. The bandwidth was set to a relatively narrow value resulting in very strong correlation, so it was expected for this result to be somewhat inaccurate.

The autocorrelation function in Figure 118 shows that there is correlation between the samples. The correlation times were measured as 10.9 samples for $\tau_{1/2}$ and 15.8 samples for $\tau_{1/e}$. The signal was a maximum at 0 samples and was completely decorrelated after approximately 100 samples. These results are very similar to the results of the underlying correlated Gaussian sequence shown in Figure 33, discussed in Section 3.4.2. Figure 119 shows that the phase of the correlated Weibull sequence is uniform between $-\pi$ and π.

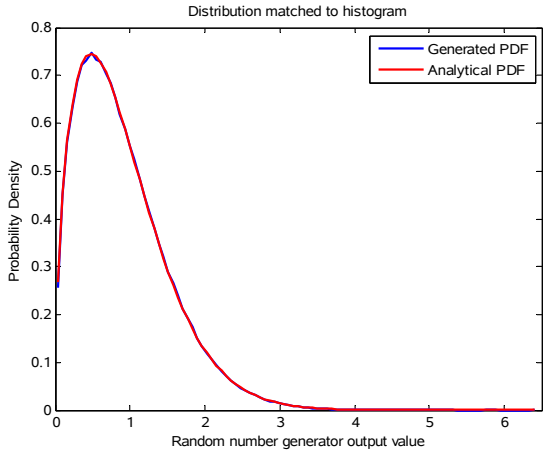

Figure 116: Normalised histogram of the time correlated Weibull random variable.

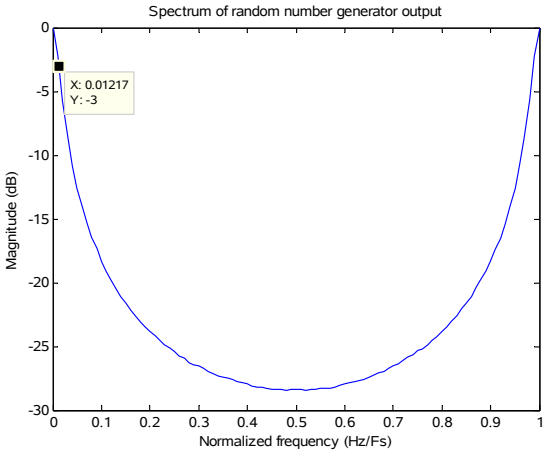

Figure 117: Frequency spectrum of the correlated Weibull random variable.

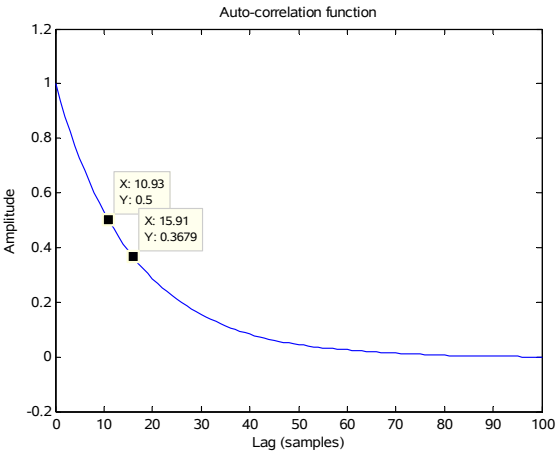

Figure 118: The autocorrelation function of the random variable.

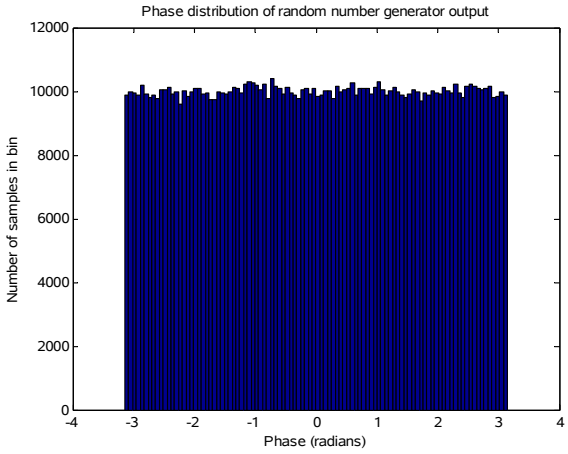

Figure 119: The phase of the correlated Weibull random variable.

The higher order moments were recorded for the correlated Weibull experiment. The higher order moments are tabulated in Table 16. The percentage error for the correlated Weibull distribution is larger than that of the uncorrelated case. The higher the order the larger the percentage error became, which hints that there are inaccuracies in the tails of the distribution, but not to the extent that was seen with the Log-Normal distribution.

Figure 120 shows the recorded moments plotted against the analytical ones. The figure shows that the measured 6th order moment is slightly smaller than the analytically calculated 6th order moment, which confirms that the samples in the tail are less than they should be. There is also a deviation from the analytical distribution from the 2nd order moment onward which indicates that the body of the distribution showed problems.

Table 16: Percentage error of the higher order moments for the correlated Weibull distribution.

Moment order	1st	2nd	3rd	4th	5th	6th
Percentage error	-0.2814%	-0.8312%	-1.6515%	-2.7721%	-4.2244%	-6.0263%

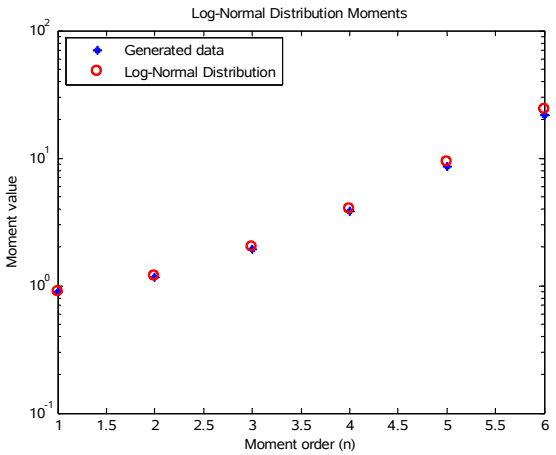

Figure 120: Measured moments from the generated correlated Weibull distribution.

Repeating the experiment 30 times and performing a χ^2 test provides a measure of how well the PDF of the generated Weibull distribution fits to the required Weibull PDF. The χ^2 test was rejected for each of the 30 experiments. The χ^2 mean rejection threshold was 102.2. The mean test outcome was 239.8 with 99% confidence band from 210.0 to 269.6. This shows that the according to the χ^2 test the generated Weibull PDF is not close enough to the analytical Weibull PDF to the specified level of confidence.

The performance of the correlated Weibull distribution was only recorded at a single bandwidth setting for this experiment. The next experiment will investigate the performance across a wide range of bandwidth settings.

3.7.3 Experiment 3: Trends with varying correlation

From the previous experiment it was observed that there is a degradation in performance between the uncorrelated and the correlated Weibull random sequence. This degradation was observed in the accuracy of the PDF and the accuracy of the higher order moments. This experiment investigates the trends in degradation in performance, by varying how strongly correlated the generated samples are. The desired bandwidth is stepped from 0.01 Fs in steps of 0.01 to 0.25 Fs, with all other variables kept constant.

Figure 121 shows the difference between the set -3 dB bandwidths and the actual measured -3 dB bandwidths. Figure 122 shows a zoomed in version of Figure 121 for narrow bandwidths. It is clearly visible that for the narrow bandwidths there is about an absolute error of 0.002, which translates to an error of 20% between the narrowest -3 dB bandwidth value and the narrowest desired -3 dB bandwidth value. As the -3 dB bandwidth becomes wider, the absolute error reduces. At the larger -3 dB bandwidths the mean of the measured -3 dB bandwidth has an absolute error of less than 0.001, which translates to an error of 0.1%. The 99% confidence intervals for the mean at the large -3 dB includes the desired -3 dB bandwidth values. The generated samples thus have the required bandwidth at wide bandwidth settings with 99% confidence.

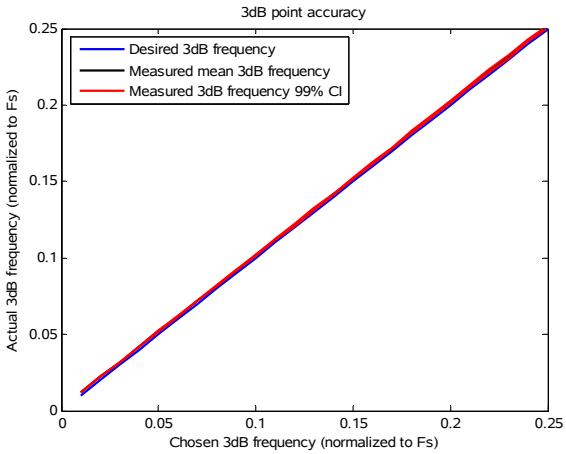

Figure 121: Measured -3 dB bandwidth of the correlated Rayleigh random number.

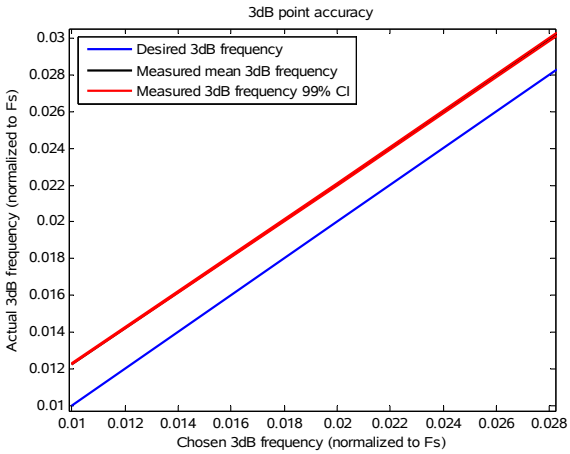

Figure 122: A zoomed in version of the graph of Figure 121 showing small bandwidth settings.

Figure 123 shows the result of χ^2 tests, averaging over 100 experiments for each value of -3 dB bandwidth. Figure 124 shows the number of χ^2 tests out of a 100 that were not rejected for each of the -3 dB bandwidth settings. From Figures 123 and 124 it can be seen that the χ^2 test rejects the majority of the generated Weibull PDFs for all correlation settings. The χ^2 test stays flatter for wider bandwidths, and starts to deteriorate with increasing rate as the bandwidth narrows. This result is very similar to the results observed for the underlying Gaussian distribution, as well as the previously discussed Rayleigh and Log-Normal distributed samples.

Figure 125 shows the KL divergence value with varying degrees of correlation. There is a trend of worsening performance as the correlation in the samples are increased. There appears to be no increase in the variance of the results obtained. The largest outlier worsens as the correlation is increased indicated by the red shot dashed line in the plot. The average result is roughly the same to that of the similar experiment performed on the correlated Gaussian distribution from Section 3.4.4, but as the correlation increases the Weibull distribution performs better than that of the Gaussian distribution. Figure 126 shows the number of samples in the tail of the Weibull distribution,

summed to give a value out of 100. There is no clear trend, and the number of samples in the tails appear to be constant irrespective of the normalised bandwidth setting. There is however an increase in variability of the result as seen by the variance shown on the plot by the blue lines, and the increase in maximum and minimum outliers shown on the plot by the red lines at the top and the bottom of the graph respectively.

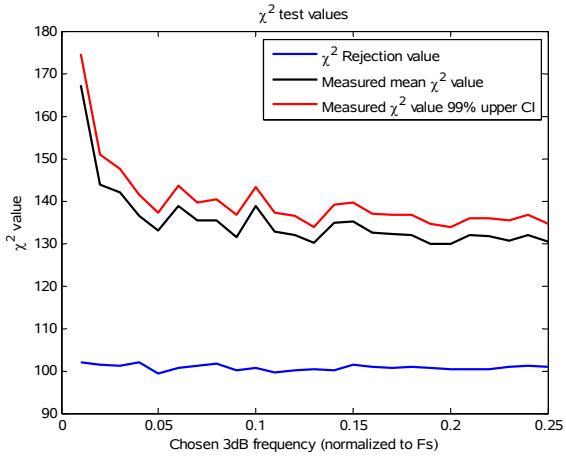

Figure 123: Plot of the outcome from a chi squared test.

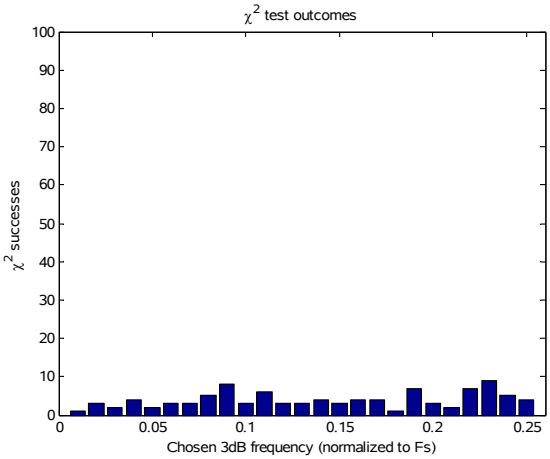

Figure 124: Number of chi squared tests that were unable to reject the PDF.

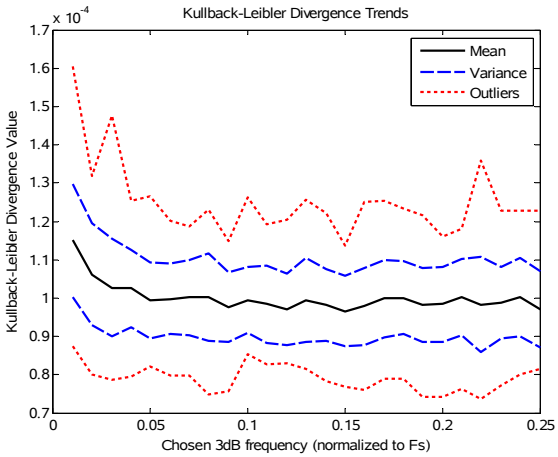

Figure 125: Kullback-Leibler divergence value.

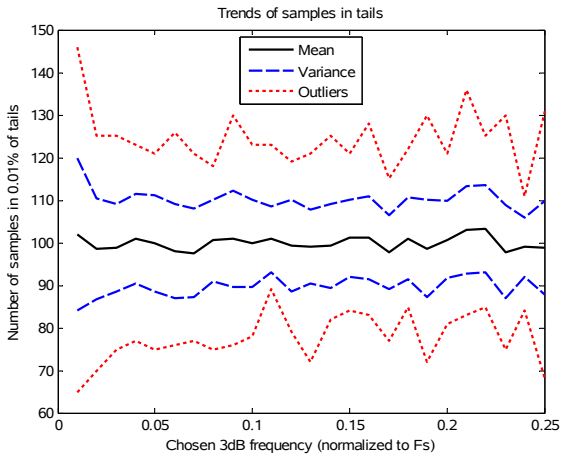

Figure 126: Number of samples in 0.01% of the tail.

Figure 127 shows the results from of the percentage error between the measured higher order moments and the calculated ones. All size the measured moments have the same trend with an error varying around zero percent irrespective of the bandwidth setting. The larger the order of the moment the larger the percentage error.

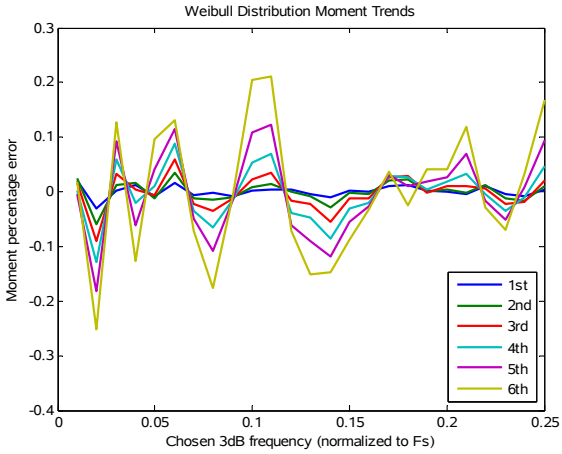

Figure 127: Percentage error of the measured moments from the generated correlated Weibull distribution.

This experiment has shown that there is a deviation from the ideal Weibull PDF for all normalised bandwidth settings. This deviation is related to the deviation trend observed in the underlying Gaussian random sequences used in the generation of the Weibull random sequence. Both the χ^2 test and the KL divergence value showed poorer performance for smaller bandwidth settings. The correlation properties were within the required bounds and showed no sign of problems. Next the performance of the Weibull generation algorithm will be analysed with the variation of the shape parameter.

3.7.4 Experiment 4: Trend with varying Weibull shape parameter

This experiment was performed with two independent sequences of Gaussian random numbers as described in the previous two experiments. For this experiment the value of the Weibull shape parameter k was varied between 1 and 5. This test was performed at a small bandwidth setting of 0.025 Fs. One million correlated samples were generated for each of the Gaussian distributed inputs to the algorithm of Figure 107, to give one million correlated Weibull distributed samples.

Figure 128 shows the result of the -3 dB bandwidth of the generated Weibull sequence as a function of the shape parameter k. The error between the desired -3 dB bandwidth and the measured bandwidth starts off at about 12% for $k=1$ and gradually decreases to 7% for $k=2.5$ where after the error increases slightly to 8% at $k=5$. Figure 129 and Figure 130 shows the results from the χ^2 tests. On average not a single value for k had a successful average χ^2 test outcome. The histogram in Figure 130 shows that out of 100 the less than 30 Weibull PDFs were rejected by the χ^2 test for most values of k. The only exception to this was between $k=1$ and $k=1.8$ where almost all the generated Weibull PDFs were rejected. The Weibull random sequence generator PDF

performed worse at $k=1.2$ and $k=1.4$, only improving slightly at $k=1.6$. For $k=1$ and k larger than 1.8 the performance was better, and roughly the same. The degradation in performance between $k=1$ and $k=2$ could be attributed to the longer tails of the Weibull distribution in this range. This suggests that the method to group samples in the tails of the distribution to minimise the effect of sparse sampling in this region has not worked as effectively as was hoped.

Figure 131 shows the KL divergence value with varying values of the Weibull distribution scale parameter. There is no trend in the performance as the shape parameter k is varied. The largest outlier, smallest outlier and variance show no observable trend in the KL divergence value. Figure 132 shows the number of samples in the tail of the Weibull distribution, summed to give a value out of 100. There is again no trend, and the number of samples in the tails appear to be constant irrespective scale parameter setting. Unlike for the previous experiment where the correlation was varied, there is no change in variability of the result as seen by the variance shown on the plot by the blue lines.

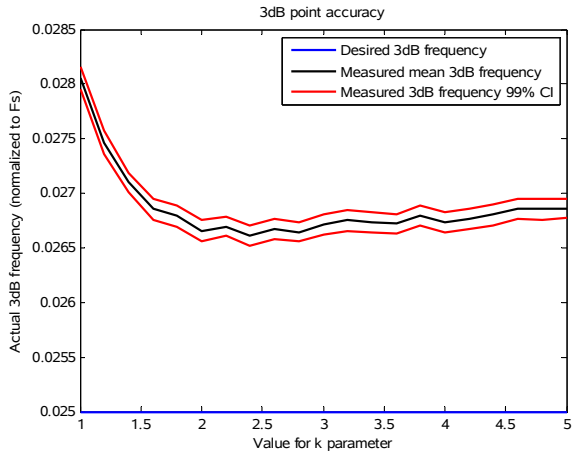

Figure 128: Measured -3 dB bandwidth of the correlated Weibull random number.

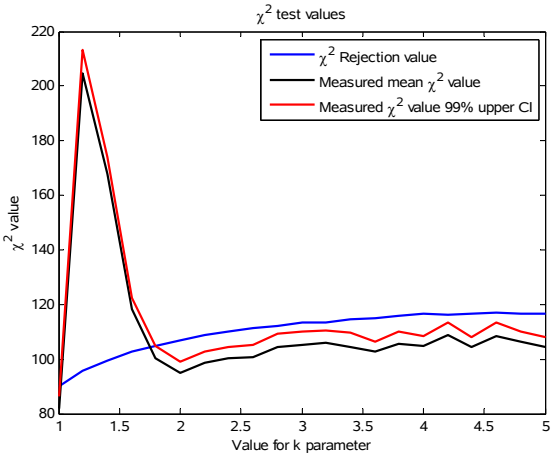

Figure 129: Plot of the outcome from a chi squared test.

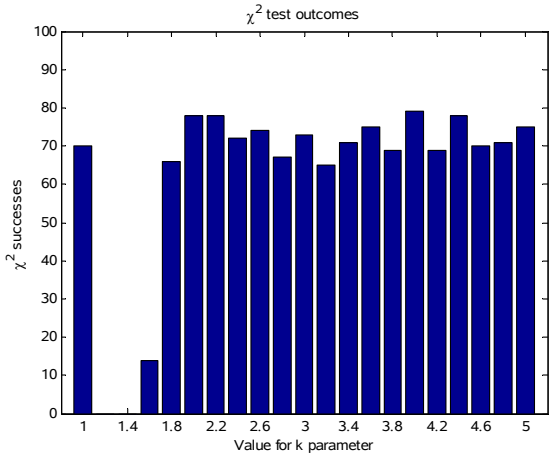

Figure 130: Number of chi squared tests that were unable to reject the PDF.

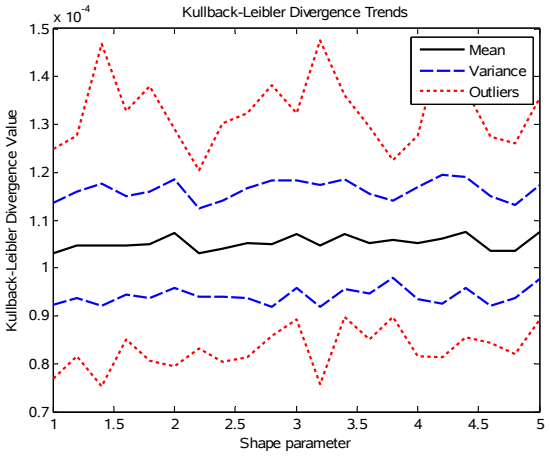

Figure 131: Kullback-Leibler divergence value.

The percentage error of the higher order moments is shown in Figure 133. The higher order moments show a trend of improvement with an increase in the shape parameter k. The improvement can mostly be noticed in the variation of the result. The 6th order moment always showed the largest error.

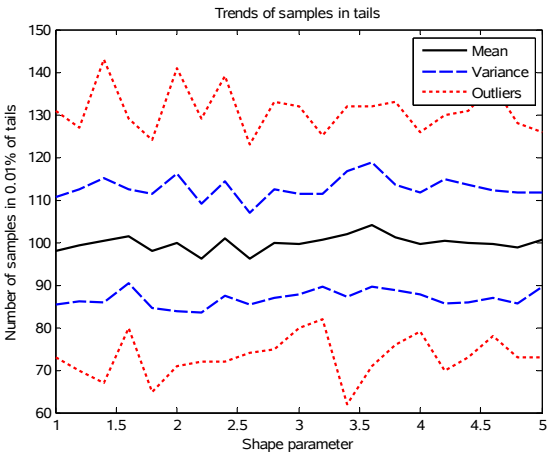

Figure 132: Number of samples in 0.01% of the tail.

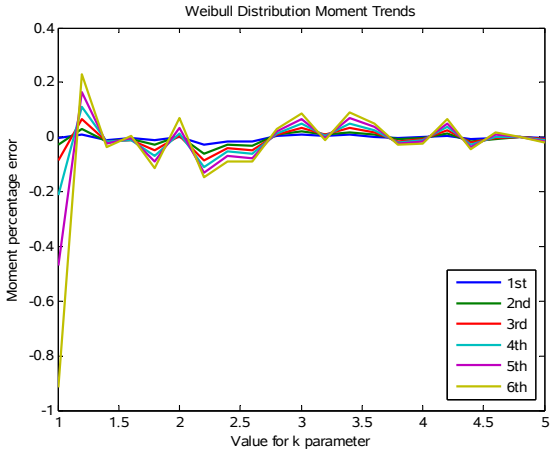

Figure 133: Percentage error of the measured moments from the generated correlated Weibull distribution.

From these results it can be seen that the Weibull distributed signal does not performs well. Although using this method to generate Weibull distributed numbers is less complex than the method used to generate Log-Normal distributed numbers, the method used for the Log-Normal distributed numbers performed better.

3.8 Conclusion

Three distributions commonly used with clutter modelling were investigated for the simulation of ground clutter. The three distributions were the Rayleigh, Log-Normal and Weibull distributions. The Gaussian distribution formed the basis for all the methods used to simulate the more complex distributions. The Gaussian distribution is unique in the sense that correlation can be added to the samples within the distribution by filtering the sequence while the resultant distribution remains Gaussian. The correlated Gaussian samples can then be manipulated to generate the required distributions for clutter simulation.

The Rayleigh distribution is the simplest to simulate, by creating a complex signal with a Gaussian distributed I and an independent Gaussian distributed Q channel. The methods to simulate the Log-Normal and Weibull distributions were more complicated than that of the Rayleigh. Both methods made use of an MNLT to transform Gaussian distributions to those of the required Log-Normal or Weibull distributions.

The PDF of the Gaussian distribution was verified by means of the χ^2 test, and the KL divergence value. The distribution of the generated uncorrelated samples were unable to be rejected by the χ^2 test as Gaussian. Multiple χ^2 tests were performed in conjunction with confidence intervals, and a confidence factor of more than 99% was obtained. There was thus a high confidence in the result that the generated samples matched the analytical Gaussian PDF and that the error between the generated PDF and the analytical PDF was negligibly small. The KL divergence produced a very small value, an indication that the generated PDF is close to that of a Gaussian. All the following experiments used the KL divergence value as the basis for further comparisons. The uncorrelated Gaussian samples showed a white frequency spectrum and an impulse in the auto correlation function as expected from signal theory. The measured moments also matched the theoretical ones.

By adding correlation to the Gaussian samples, the frequency spectrum was shaped to have a specific -3 dB bandwidth. The autocorrelation function widened from an impulse to a function with gradual decay. For the first test case the samples were weakly correlated. The added correlation degraded the performance of the PDF somewhat. The KL divergence value was slightly larger than for the uncorrelated case and the χ^2 test showed that there was still a 99% confidence that the generated distribution was Gaussian. An increase in correlation narrowed the frequency spectrum and widened the autocorrelation function even more, as expected, resulting in strongly correlated samples. The moments remained unchanged, which is also consistent with the expected behaviour. The increase in correlation exhibited poorer performance for the accuracy of the bandwidth and the PDF shape. The KL divergence value was larger than for the weakly correlated samples. The χ^2 test no longer had a 99% confidence level that the generated distribution was Gaussian. This variation in performance with variation in correlation strength indicated that there might be a negative trend in PDF quality as the correlation is increased.

Varying the correlation by changing the bandwidth, allowed the performance trends of the PDF and bandwidth to be recorded for the Gaussian distribution. The measured bandwidth matched closely that of the desired setting except for narrow bandwidths. The performance for the PDF measured with the χ^2 test was relatively unchanged across most of the bandwidth settings, but degraded for smaller bandwidths. The KL divergence value, and the higher order moments showed the same trend as the χ^2 test, except the degradation in performance was more gradual as the bandwidth was reduced. Overall the Gaussian distribution has been found to perform well. The knowledge obtained regarding its performance allowed for an informed investigation into the more complex Rayleigh, Log-Normal and Weibull distributions.

For all the tested distributions, when correlation was added to the generated PDF, an increase in error relative to the analytical PDF was observed, as compared to the uncorrelated case. By increasing the correlation (reducing the bandwidth) there was an increase in error observed for the generated PDF from the analytical PDF. This was true for the underlying Gaussian as discussed in the previous paragraph, but also true for the generated Rayleigh, Log-Normal and Weibull distributions. All the generated distributions degraded gracefully, by slowly deviating from the analytical PDF as the correlation was increased. The χ^2 test showed that the generated Rayleigh distribution was the closest fit to its analytical distribution followed by the Log-Normal and then the Weibull distribution. The χ^2 test failed to reject on average 90% of the Rayleigh distribution, 50% of the Log-Normal and only 10% of the Weibull distribution. The KL divergence values showed nearly identical results for all three distributions.

The -3 dB bandwidth of the frequency spectrum for the generated PDFs was controlled by the -3 dB bandwidth of the underlying correlated Gaussian random sequence as this is a property of an MNLT. The measured -3 dB bandwidths were in general very close to the desired -3 dB bandwidth values. As discussed above, the Gaussian distribution had a small error between the measured -3 dB bandwidth and the desired -3 dB bandwidth for wide bandwidth settings and the error became larger as the desired -3 dB bandwidth decreased. This same trend was visible for the Rayleigh Log-Normal and Weibull distributions. The autocorrelation function was well behaved for the simulated Rayleigh, Log-Normal and Weibull clutter, and exhibited similar results to those of the underlying Gaussian random number generator.

By varying the input parameters, α for the Rayleigh distribution, σ for the Log-Normal distribution, and k for the Weibull distribution, trends in the χ^2 test outcomes could be observed as a function of the controlling parameters of the individual PDFs. For each of these tests the correlation was held constant by setting the bandwidth of the generated random numbers to $Fs = 0.025$. The Rayleigh distribution did not show any trend in χ^2 value or KL divergence value when α was varied. The χ^2 test failed to reject the generated Rayleigh PDF on average. This is a positive result as it means the performance of the Rayleigh distributed samples are independent of the scale parameter. The Log-Normal distribution behaved well for values of σ below 0.5, and very poorly for values of σ larger than 0.5. The change between acceptable and poor performance was near instantaneous. The KL divergence value showed a slight decease in size contradicting the results of the χ^2 test. The apparent improvement in performance shown by the KL divergence value might be due to the fact that the KL divergence does not count empty histogram bins. Both

tests indicated that the performance of the Log-Normal distributed samples were not independent of its scale parameter. This could be due to the fact that the "scale parameter" of the Log-Normal distribution is not an accurate description since it also alters the shape of the distribution, while the scale parameter for the Rayleigh distribution only affects the scale. The Weibull distribution was not rejected by the χ^2 test for almost every value of k. The outcomes of the χ^2 test were constant for most values of k except $k=1.4$ where the observed performance was much worse. The performance of the Weibull distributed samples was fairly flat except for the outlier at $k=1.4$ which means the performance observed was relatively unchanged by the shape parameter. The outlier at $k=1.4$ can be attributed to the longer tails exhibited by the Weibull distribution shape for this setting.

At no point during any of the experiments did measuring the number of samples in the tail and comparing it to the expected number of samples for that section of the tail, show a trend indicative of a change in performance of the generated distributions. Even distributions with extremely long tails such as the Log-Normal, measured over the very end of the tails (last 10 samples or 0.001% of all generated samples), showed no deviation from the expected number of samples.

The moments did not show any clear trend in performance of the generated distribution when the average values were used. The only real difference was that the results varied more drastically when the samples were more strong correlated.

A uniform phase between -180° to 180° is a requirement for radar clutter simulation at high RF frequencies. Each of the respective methods tested for the Rayleigh distributed, Log-Normal distributed or the Weibull distributed clutter produced a uniform phase from -180° to 180° which meets the requirement for the phase distribution.

The poor performance observed for narrow bandwidth settings for the Gaussian, Rayleigh, Log-Normal and Weibull distributions could possibly be caused by the behaviour of a correlated signal. Since correlated samples are limited by the size of the amplitude step which can be taken from one sample to the next, this could cause an even lower probability of samples to be generated in the tails, similarly to the case where too few samples of a given distribution are generated. This lack of samples in the tails becomes more severe as the correlation is increased. An investigation into the behaviour of these generated distributions by varying the number of samples when sequences with narrow bandwidths are generated, will be worthwhile. Additionally an experiment to determine if a

plateau is reached in the accuracy of the generated distributions versus the number of samples will also be beneficial. This approach could better answer the question as to whether a distribution is generated poorly, but it will require more computational resources.

Of the methods identified as suitable for implementation on an FPGA both the Log-Normal and Weibull distributions exhibited good performance. Unfortunately for both the Log-Normal distribution and the Weibull distribution, the complexity of the method will be the main drawback to FPGA implementation. Both these methods had many processing steps, some of them mathematically intensive, which is also not ideal. A different approach is thus required, starting from the ground up, if the eventual implementation is to be on an FPGA based platform. What are the difficulties of simulating a complex signal with a certain phase and magnitude, by starting from exactly that? Could this lead to an approach that has low complexity and is more suited to implementation on an FPGA? An investigation into a method that is more optimal for implementation on an FPGA was developed next to test its performance and to place the strengths of the methods from literature in perspective.

4 SIMULATION OF CLUTTER, A NUMERICAL APPROACH

Since the methods to simulate clutter that were identified in the literature were found to be lacking simplicity, a different approach was developed and investigated to identify the weaknesses of the techniques, and determine whether these techniques can be improved upon. The approach starts from the radar, where the distribution of the detected energy in the environment is the result of taking the magnitude of a complex signal. This complex signal also has a uniform phase, and a specific Doppler bandwidth depending on the motion of both the platform and the environment.

A uniform PDF is evenly distributed between 0 and 1. The x-axis of a CDF for any distribution can span from $-\infty$ to $+\infty$. Since the CDF is the integral of the distribution, it can never have a values less than 0, or have values of greater than 1. This flows from the definition of a PDF, where the surface area underneath the curve is equal to unity. A CDF is also a monotonically increasing function. These properties can be exploited to generate a specified PDF using a technique called CDF inversion. CDF inversion is a technique that can generate a random number with any distribution that can be described numerically [43]. The inverse of the CDF lies between 0 and 1, and is also a one-to-one function in the sense that there is only one value for y for each value of x. When random values from a uniform PDF are transformed through the inverse cumulative distribution function (ICDF) the resulting distribution is that of the PDF which corresponds to the ICDF. This property is very attractive since it is relatively easy to generate a uniformly distributed random variable digitally. The inverse is also true for the CDF. When values of a certain distribution is transformed through that specific distribution's CDF, the resulting PDF is uniformly distributed between 0 and 1. These two directions to the same concept (as illustrated in Figure 134) are used with this method to generate a correlated clutter variable.

Generating a correlated Gaussian random variable has been discussed in Section 3.2. This method is used in conjunction with the Gaussian CDF to produce a correlated uniform random variable. The correlated uniform random variable is generated by the ICDF technique to form a uniformly distributed RV, which is then used as the phase ϕ of the random number sequence. Another independent correlated uniform random variable is generated with the same method in parallel with the first. This correlated uniform random variable is passed through an inverted CDF to obtain the magnitude R which has the desired amplitude PDF. Together the magnitude and phase produces a

correlated random number sequence. The complex envelope representation is converted into the real and imaginary parts of a signal using equations (49) and (50). Figure 135 illustrates this process.

$$x_{\Re} = R.cos(\phi) \tag{49}$$

$$x_{\Im} = R.sin(\phi) \tag{50}$$

The method described here will be referred to as the CDF transformation method to distinguish it from the CDF inversion method. The CDF transformation method will be investigated to determine its performance relative to the methods implemented from literature in Section 3.

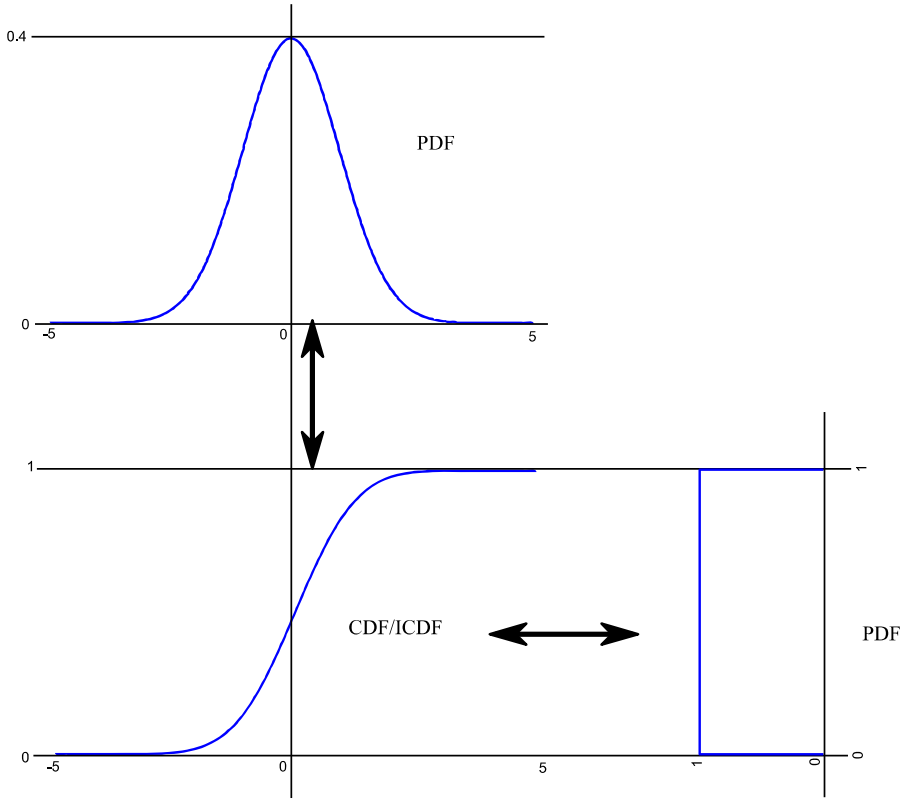

Figure 134: An illustration showing the relationships between a PDF, CDF and ICDF.

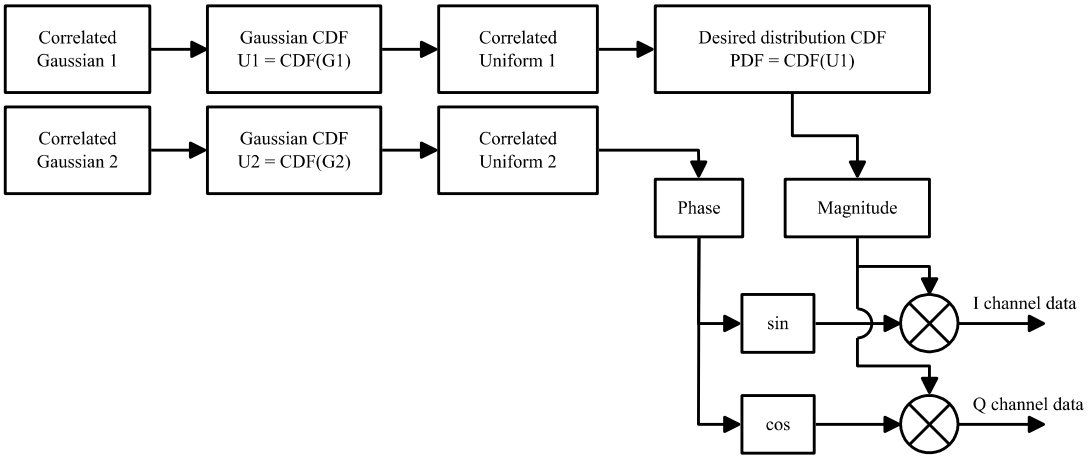

Figure 135: Block diagram of the clutter generation method with the CDF transformation method.

4.1 Correlated uniformly distributed phase

This section describes the steps that are independent of the chosen distribution. The most important of these steps is the generation of uniformly distributed phase variables. This is achieved by generating a correlated uniformly distributed random sequence. The correlated uniform random variable is created by generating a correlated Gaussian distribution, and then passing the generated values through the equation of the Gaussian CDF given in equation (51). Each sample of the Gaussian distribution is used as the input x to create a new value y. A histogram of the transformed values has the required uniform distribution. The parameters μ and σ of the transformation function are the mean and standard deviation of the generated samples. Ideally these should be constant and independent of the bandwidth of the correlated Gaussian distribution. The method used to generate the correlated Gaussian random numbers is the same method of single pole filtering described in Section 3.2 and analysed in Section 3.4.2.

$$y=\frac{1}{2}\left[1+erf\left(\frac{x-\mu}{\sqrt{2\sigma^2}}\right)\right] \tag{51}$$

Figure 136 shows the resultant distribution after passing the Gaussian distributed samples through equation (51), shifting with -0.5, and multiplying by 2π to obtain a uniform distribution in the range of $-\pi$ to π. This uniform distribution acts as the phase for the correlated clutter variables generated with equations (49) and (50). Figure 137 shows the frequency spectrum of the correlated uniformly distributed random variable plotted over that of the correlated Gaussian distributed random variable. The -3 dB bandwidth setting was chosen to be 0.01 Fs, after the transformation from Gaussian to uniform the -3 dB bandwidth was 0.0146 Fs. There difference observed in the -3 dB bandwidth between the Gaussian and uniform distributed random variables were close enough. The correlated Gaussian distribution used as input for the transformation had shallow (approximately 1 dB) notches in the frequency spectrum at 0.5 Fs (shown in Figure 137 as well as Figure 32 on page 39 of Section 3.4.3). After the transformation the notch at 0.5 Fs is still present. The only noticeable difference is the slightly elevated noise floor of the frequency spectrum, which is related to the bandwidth. This aspect of the figure suggests that when the variable is transformed from a Gaussian distributed to a uniform distributed sequence that some decorrelation occurs. This decorrelation in the sequence of random samples is expected. The methods from literature (Section 3) aim to preserve the correlation after transformation, and no such attempt has been made with the algorithm discussed here.

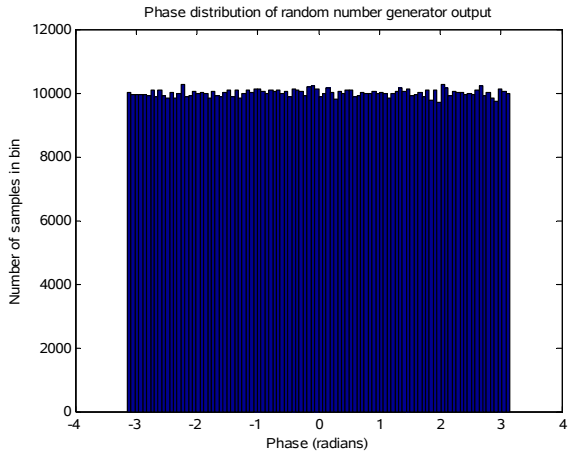

Figure 136: Distribution of the correlated uniform random variable.

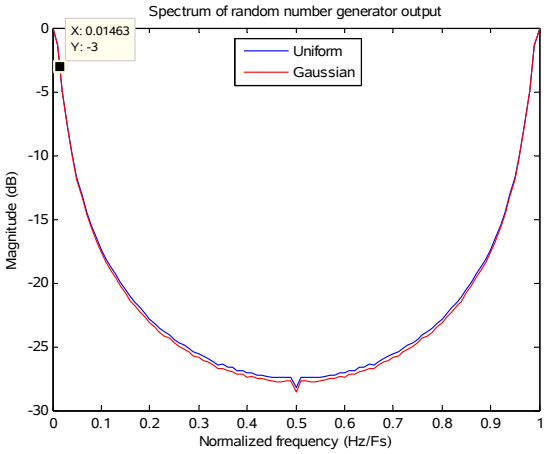

Figure 137: Frequency spectrum of the correlated uniform and Gaussian random variable.

Figure 138 shows a comparison of the measured -3 dB bandwidths of both the uniform and Gaussian distributed sequences, plotted against the desired -3 dB bandwidths used as input. The normalised bandwidth was varied from 0.01 Fs to 0.25 Fs. A total 30 measurements were made and averaged for each step in Fs in the plots. Figure 139 shows the error between the set -3 dB bandwidths and the actual measured -3 dB bandwidths for the uniform distributed sequence and the Gaussian distributed sequence it was transformed from. The Gaussian distributed sequence initially started with a relatively large error but as the bandwidth increases, it tends to the desired -3 dB bandwidth. The uniformly distributed sequence starts off the same, but does not tend to go to the input -3 dB bandwidth as the Gaussian distributed sequence did, and instead the error increases gradually. There is thus some error introduced when the random sequence is transformed from the Gaussian distribution to the uniform distribution, but this error is negligibly small.

Figure 140 shows the autocorrelation functions of the uniform distributed sequence and the Gaussian distributed sequence. The autocorrelation function shows that there is a similar correlation between the two sequences, and that it dies out to approximately zero after 90 samples. The decorrelation values for the uniform distributed sequence was 10.5 samples for $\tau_{1/2}$, and 15.4 samples for $\tau_{1/e}$. The values for the -3 dB bandwidth, $\tau_{1/2}$ and $\tau_{1/e}$ are comparable to those measured before the transformation was applied as seen in Figures 32 and 33 in Section 3.4.2.

The same transformation without the multiplication and shifting as done here for the phase is used to obtain a uniform distribution in the range of 0 to 1 as an intermediary step between the correlated Gaussian and the distribution for the magnitude of the clutter. The Rayleigh, Log-Normal and Weibull distributions have been investigated as magnitudes for direct comparison to the results obtained in Section 3.

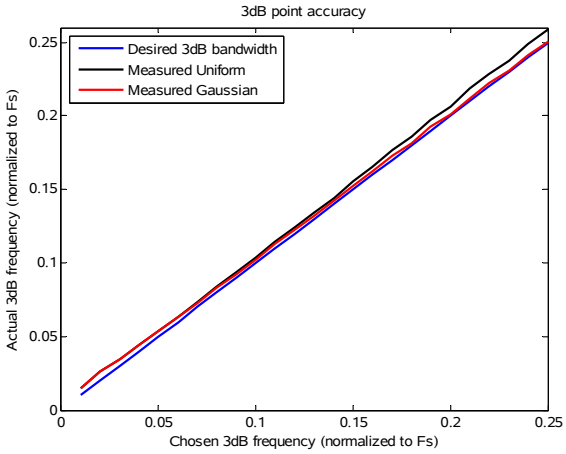

Figure 138: Measured -3 dB bandwidth of the correlated uniform and Gaussian random number.

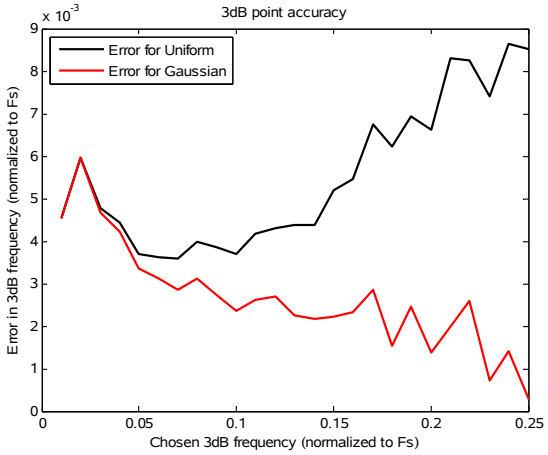

Figure 139: Comparison of error in the measured -3 dB bandwidth.

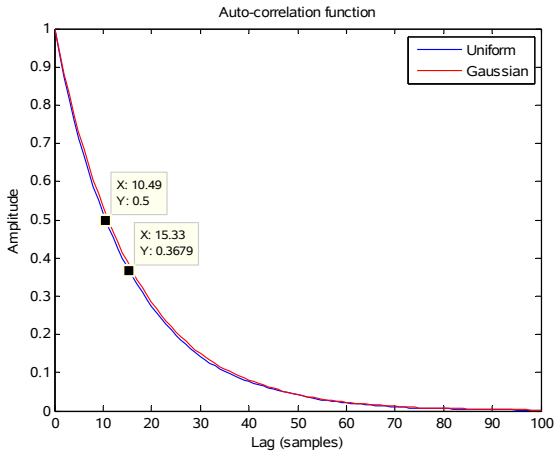

Figure 140: The autocorrelation function of the correlated uniform and Gaussian random variable.

4.2 Rayleigh distributed clutter

The Rayleigh distribution was generated with the CDF transformation method as the magnitude of the clutter. An uncorrelated uniform variable is generated with the range 0 to 1. The resulting random values are then passed through the inverted CDF of the Rayleigh distribution. The CDF of the Rayleigh distribution is shown in equation (52) and the inverse CDF in equation (53).

$$y = 1 - e^{\frac{-x^2}{2\alpha^2}} \tag{52}$$

$$x = \sqrt{-2\alpha^2 \ln(1-y)} \tag{53}$$

For brevity only the results that capture the overall performance, or highlight certain differences between this method and the previous methods from literature have been analysed.

4.2.1 Random sample generation

This experiment serves to show the performance of a single simulation run. The PDF, moments, frequency spectrum and autocorrelation function are analysed to ensure the correct behaviour is expressed. This was done to confirm that the CDF transformation method produces a PDF that has the correct shape, moments that agree with the theoretically calculated ones, a smooth monotonically decreasing frequency spectrum without large deviations from the theoretical shape, and an autocorrelation function that shows decorrelation to the extent that the result decays to zero. This experiment is an initial check before the more rigorous experiments are performed.

For this experiment a Rayleigh random number sequence has been generated, with a value for α set to 1 and the bandwidth set to 0.025 Fs. The CDF transformation method discussed at the start of this section was used to generate a correlated phase, and a correlated magnitude to the shape of a Rayleigh PDF. A total of one million correlated samples were generated.

Figure 141 shows the analytical PDF plotted against the histogram of the generated samples. There is a close fit between the analytical PDF and the measured samples. Figure 142 shows the measured moments obtained from the generated data, compared to the theoretical moments for the Rayleigh distribution. There is a close fit for the moments of lower order, and a small but negligible deviation is noticeable for the 5th and the 6th order moments.

Performing the KL divergence test on the correlated Rayleigh random samples averaged over 30 experiments resulted in a value of 0.00011265. This value is 2.11% less than that of the correlated Rayleigh distribution generated in Section 3.5.2. The main reason for the improved performance could be that the correlation for the original Rayleigh distributed samples in Section 3.5.2 was set to 0.01 Fs, where for this experiment with the CDF transformation method the correlation is 0.025 Fs. Previous results have shown that the performance tends to decrease as the correlation is increased.

Figure 143 shows the frequency spectrum of the correlated Rayleigh distributed samples generated with the CDF transformation method. The frequency spectrum is normalised to sampling frequency with the magnitudes taken relative to the largest frequency component. The frequency spectrum shows that there is correlation, but the measured bandwidth is somewhat larger than expected. The bandwidth was set to 0.025 Fs, but 0.04878 Fs, a value closer to twice the desired one was measured. Figure 144 shows the autocorrelation function for the generated sequence. The sequence

is a maximum at a delay of zero samples. The sequence decorrelates completely at around 30 samples, which is much less than seen in Figure 140, which is expected due to the larger bandwidth observed. The correlation times were measured as 1.48 samples for $\tau_{1/2}$ and 2.23 samples for $\tau_{1/e}$.

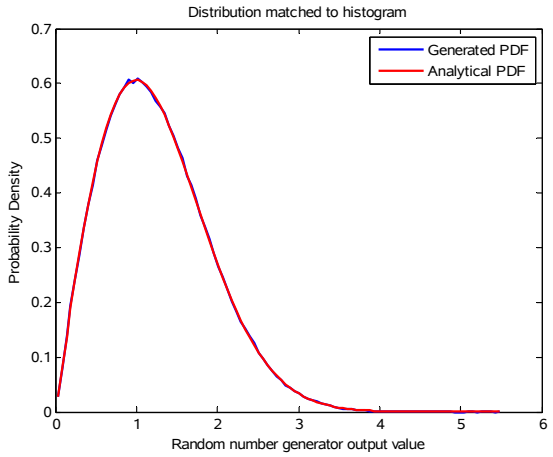

Figure 141: Normalised histogram of the time correlated Rayleigh random variable.

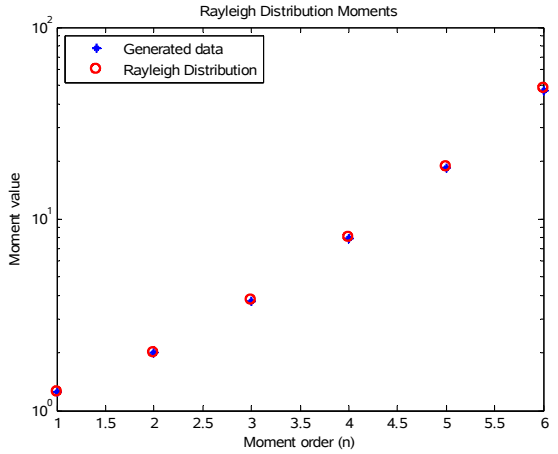

Figure 142: Measured moments from the generated correlated Rayleigh distribution.

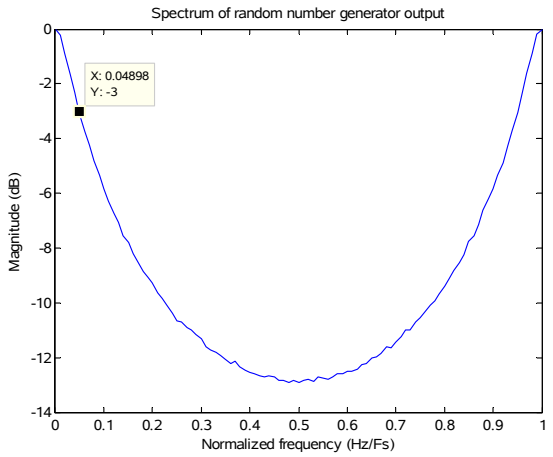

Figure 143: Frequency spectrum of the correlated Rayleigh random variable.

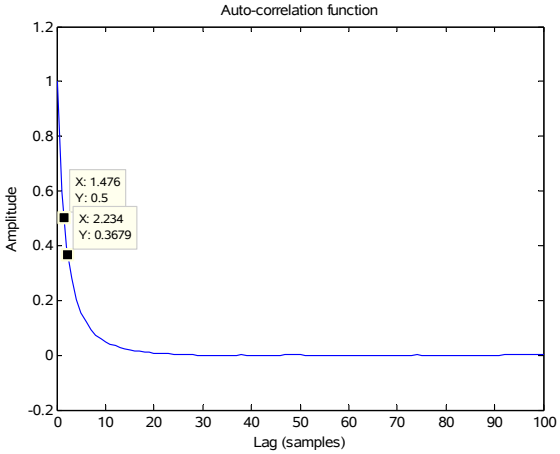

Figure 144: The autocorrelation function of the correlated random variable.

This experiment has shown that the correlated Rayleigh sequence generated with the CDF transformation method produces valid results. The Rayleigh PDF was visually accurate, and the measured moments matched that of the theoretical ones. There are some inaccuracies observed in the bandwidth of the sequence, and thus due to the duality between the time and frequency domain, the decorrelation times were less than expected. The experiments that follow will investigate this deviation from the desired bandwidth further.

4.2.2 Trend with varying correlation

This experiment investigates the trends in degradation in performance, by varying how strongly correlated the generated samples are. The desired bandwidth is stepped from 0.01 Fs to 0.25 Fs, with all other variables held constant.

From the equivalent experiment in Section 3.5.3 excellent performance was observed. The measured -3 dB bandwidth was closely matched to the desired -3 dB bandwidth, with a small error for small normalised bandwidth settings, which improved with an increase in the normalised bandwidth setting. The quality of the PDF shape was also favourable with the χ^2 test not being able to reject the PDF upward of 0.025 Fs. The percentage error of the higher order moments were less than 1% across the entire range of bandwidths, only increasing for normalised bandwidth values of less than 0.05 Fs.

Figure 145 shows the error between the set -3 dB bandwidths and the measured -3 dB bandwidths. The trend observed here is different to that observed from the MNLT methods described in Section 3, where the measured normalised bandwidth closely matched the desired normalised bandwidth. With the CDF transformation method there is a large and non-constant error between the desired normalised bandwidth and the measured normalised bandwidth. The error at 0.01 Fs is 0.112 compared to 0.002 for the previous method. The error becomes larger as the normalised bandwidth setting increases. Before the desired normalised bandwidth reaches 0.1 Fs the measured normalised bandwidth had already exceeded 0.25 Fs. We can conclude that this new method does not preserve the bandwidth of the Gaussian random number sequence used in the transformation.

Figure 146 shows the result of χ^2 tests, averaging over 100 experiments for each value of -3 dB bandwidth. It can be seen that the χ^2 test on average does not reject the generated Rayleigh PDF. The performance for this method is slightly less than that observed in the previous section. The normalised bandwidth point below which the χ^2 rejects the PDF is on average is slightly larger than that observed previously. The threshold for the χ^2 test still lies just below 110, and the asymptotic region still lies between 80 and 90.

Figure 147 shows the number of χ^2 tests out of a 100 that were not rejected for each of the -3 dB bandwidth settings. The same trend observed with the χ^2 average is also observed here. When compared to the same result from Section 3.5.3 it can be seen that the χ^2 test rejected a few more PDFs in the region above 0.05 Fs but the performance is still favourable. The point where the

PDFs rejected started to noticeably increase moved up from 0.04 Fs to 0.11 Fs. The worst performance was at the lowest normalised bandwidth setting of 0.01 Fs where only 10 PDFs were not rejected by the χ^2 test.

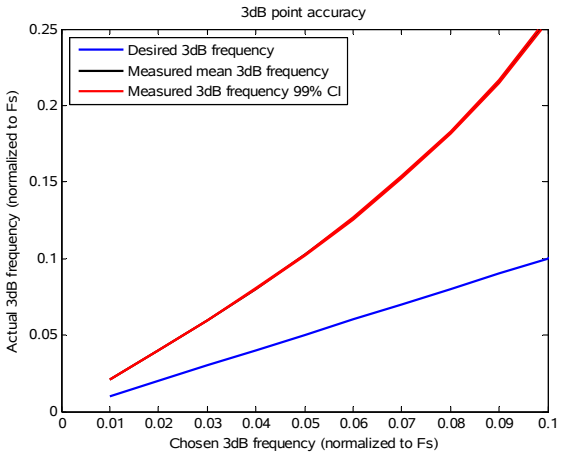

Figure 145: Measured -3 dB bandwidth of the correlated Rayleigh random number.

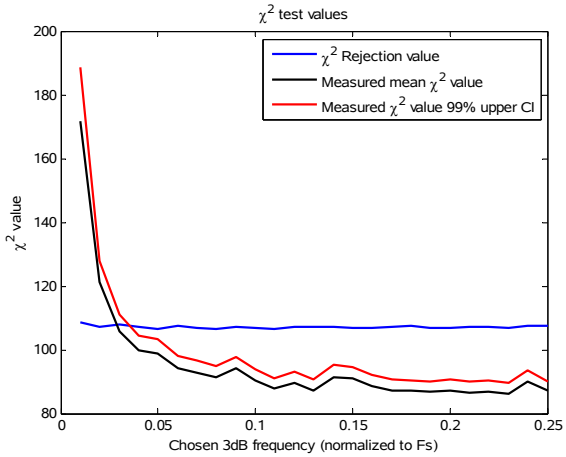

Figure 146: Plot of the outcome from a chi squared test.

Figure 148 shows the KL divergence value with varying degrees of correlation. There is a trend of worsening performance as the correlation is increased. There appears to be no increase in the variance of the results obtained except for the largest correlation at 0.01 Fs. The largest outlier worsens as the correlation is increased indicated by the red dashed line in the plot. Figure 149 shows the number of samples in the tail of the Rayleigh distribution, summed to give a value out of 100. There is no clear trend, and the number of samples in the tail appears to be constant irrespective of the normalised bandwidth setting. There is however an increase in variability of the result as seen by the variance shown on the plot by the blue lines, and the increase in maximum and minimum outliers shown on the plot by the red lines at the top and the bottom of the graph respectively.

Figure 150 shows the results from of the percentage error between the measured higher order moments and the calculated ones. The trend observed for the χ^2 in Figure 146 is not visible for the higher order moments. All six the moments have the same trend, but there is no clear variation connected with the bandwidth of the generated correlated Rayleigh samples. The larger the order of the moment the larger the percentage error, and the sooner the increase in percentage error due to the smaller bandwidth can be observed. The 6th order moment stayed within the ±0.15% error level for almost all the bandwidth settings between 0.01 Fs and 0.25 Fs. This is slightly worse than for the corresponding experiment performed in Section 3.5.3.

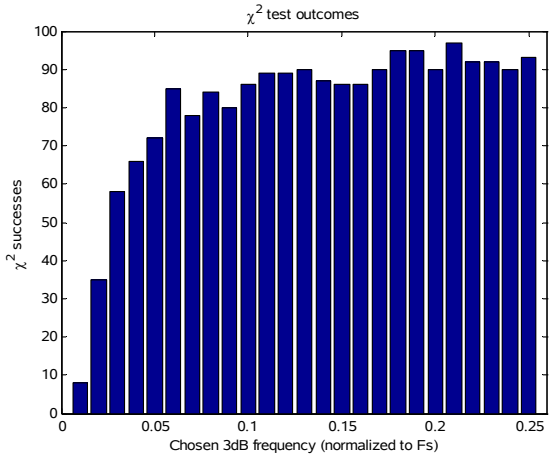

Figure 147: Number of chi squared tests that were unable to reject the PDF.

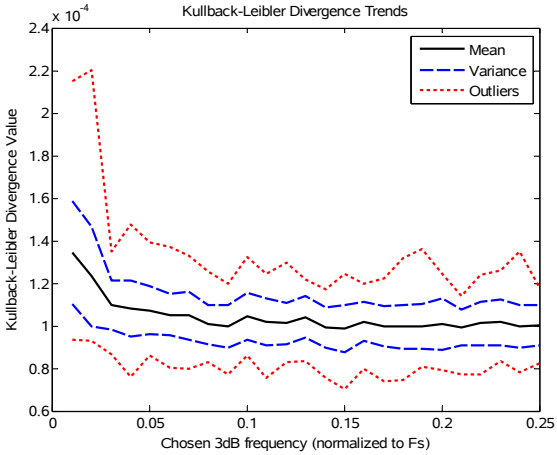

Figure 148: Kullback-Leibler divergence value.

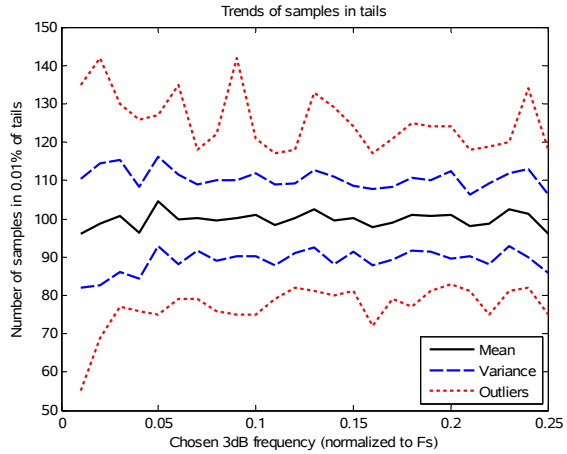

Figure 149: Number of samples in 0.01% of the tail.

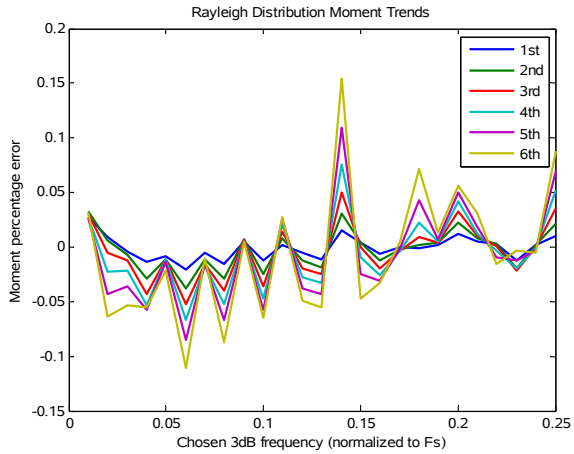

Figure 150: Percentage error of the measured moments from the generated correlated Rayleigh distribution.

This experiment has shown that except for the normalised bandwidth the performance of the CDF transformation method to generate samples with a specified PDF and uniform phase distribution performed well. The performance was slightly worse than that of the method described in Section 3.5.3 when comparing the χ^2 test, KL divergence value and the percentage error of the higher order moments. The bandwidth performance was much worse but since the relationship between the desired -3 dB normalised bandwidth and the measured -3 dB normalised bandwidth is that of a monotonic function it does not render this method completely useless.

4.2.3 Trend with varying Rayleigh scale parameter

This experiment investigates the trends in degradation in performance, by varying the Rayleigh scale parameter α. For this experiment the value of the Rayleigh scale parameter was varied between 0.5 and 3. The bandwidth of the signal was set to 0.025 Fs.

From the equivalent experiment in Section 3.5.4 excellent performance was observed. The measured -3 dB bandwidth did not show any trend for a change in the Rayleigh scale parameter α across the entire range tested. The quality of the PDF shape was also favourable with the χ^2 test not being able to reject the PDF over the same range of values for α. The percentage error of the higher order moments was less than 1% across the entire range α values.

The -3 dB accuracy is shown in Figure 151 for values of α between 0.5 and 3. The trend observed is the same as the one observed previously. The measured -3 dB bandwidth is constant across the entire range of values for the Rayleigh scale parameter α, the main difference being that the error between the desired normalised bandwidth and the measured normalised bandwidth is large. This can be expected from the results obtained in the previous experiment. Although this large error exists, this result is acceptable because it means that when the error in the bandwidth is compensated for, any value of α can be used for the Rayleigh distribution without any changes to the bandwidth.

Figure 152 shows the result of χ^2 tests, averaged over 100 experiments for each value of α. Both the 99% confidence interval and the average χ^2 test result is slightly higher than the threshold for all values of α. There is some variation which can be attributed to noise in the χ^2 test due to the variation of the random samples, but there is no trend in the results. There was also no trend observed for the previous test in Section 3.5.4 The main difference is for the shape Rayleigh PDF, which was better than with this new method.

Figure 153 shows the number of χ^2 tests out of a 100 that were not rejected for each of the values for α. The average number of tests not rejected averages to 45, compared to 75 for that of the experiment in Section 3.5.4. Similarly to the Rayleigh distribution from Section 3.5.4 there is also no clear trend visible from this result.

Figure 154 shows the KL divergence value with varying values of the Rayleigh distribution scale parameter. There is no trend in the performance as the scale parameter α is varied. The largest outlier, smallest outlier and variance show no observable trend in the KL divergence value. The average KL divergence value is slightly larger than the corresponding test performed on the Rayleigh distribution from Section 3.5.4.

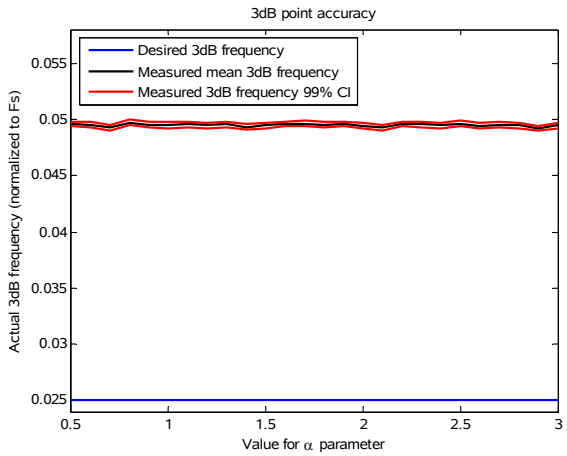

Figure 151: Measured -3 dB bandwidth of the correlated Rayleigh random number.

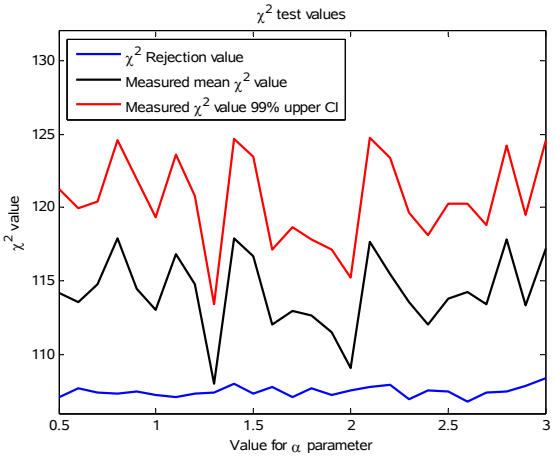

Figure 152: Plot of the outcome from a chi squared test.

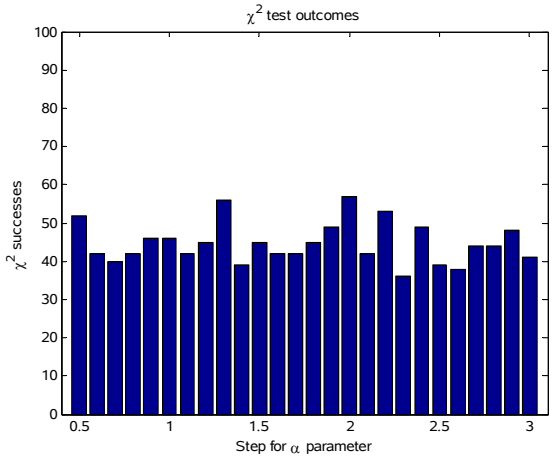

Figure 153: Number of chi squared tests that were unable to reject the PDF.

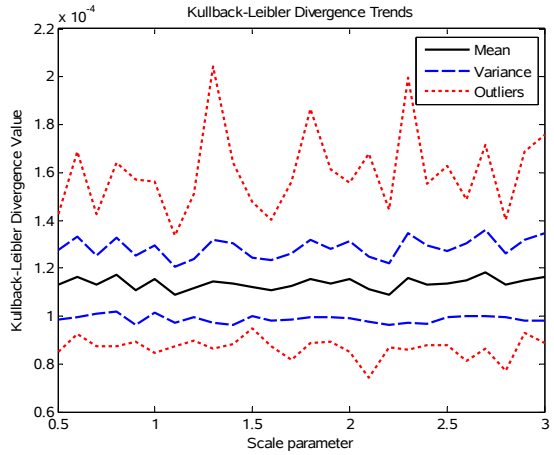

Figure 154: Kullback-Leibler divergence value.

Figure 155 shows the number of samples in the tail of the Rayleigh distribution, summed to give a value out of 100. There is again no trend, and the number of samples in the tails appear to be constant irrespective scale parameter setting. Unlike the previous experiment where the correlation was varied, there is no change in variability of the result as seen by the variance shown on the plot by the blue lines. Figure 156 shows the results from of the percentage error between the measured higher order moments and the calculated ones. Similarly to the χ^2 test, the moments reveal no trend in the performance of the generated Rayleigh distribution for all values of α tested. The 6th order moment stayed below the 0.25% error level for the entire range of values for α. Compared to the results in Section 3.5.4 there was a slight increase in the average percentage error but this increase is less than 0.05%.

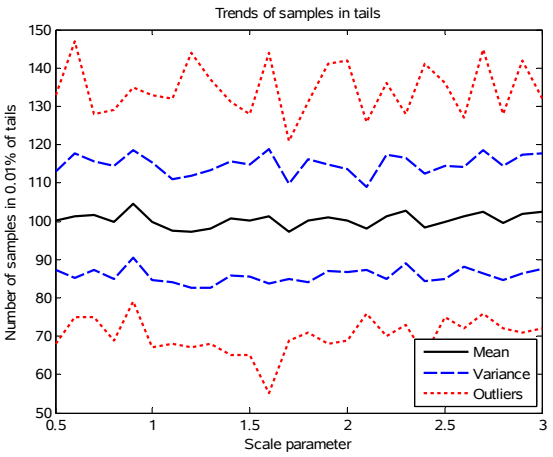

Figure 155: Number of samples in 0.01% of the tail.

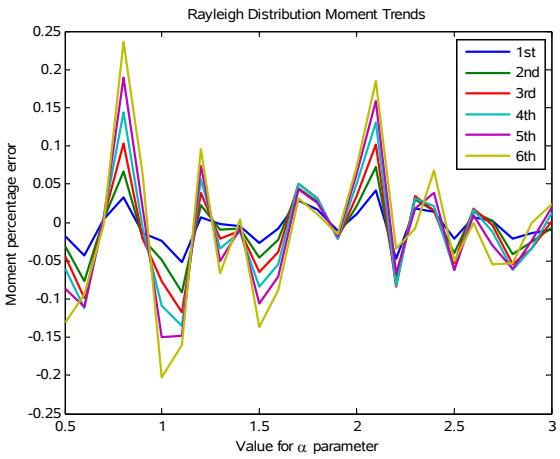

Figure 156: Percentage error of the measured moments from the generated correlated Rayleigh distribution.

The results obtained from this experiment show that there is still no measurable variation in the performance of the Rayleigh distributed correlated random sequence when the scale parameter α is varied. Some performance degradation is noticeable, but this poorer performance is aligned with the results from the previous experiment in this section. A noticeable drawback is that since the measured normalised bandwidth is much higher than the desired normalised bandwidth, the measured performance is somewhat worse than indicated here. If the desired bandwidth is set to 0.05 Fs right where the degradation in performance was noticed, then the measured -3 dB bandwidth is closer to 0.1 Fs. Thus the poor performance is at a much wider bandwidth than the results indicate, since the results are plotted against the desired bandwidth, and not the actual measured bandwidth.

4.3 Log-Normal distributed clutter

The Log-Normal distribution was generated to serve as the magnitude of the clutter. An uncorrelated uniform variable is generated in the range between 0 and 1. The resulting random values are then passed through the inverted CDF of the Log-Normal distribution. The CDF of the Log-Normal distribution is shown in equation (54) and the inverse CDF in equation (55).

$$y = \frac{1}{2} + \frac{1}{2} erf\left[\frac{\ln(x) - \mu}{\sqrt{2\sigma^2}}\right] \tag{54}$$

$$x = e^{\sqrt{2\sigma^2} erf^{-1}(2y-1) + \mu} \tag{55}$$

For brevity only the results that capture the overall performance, or highlight certain differences between this method and the previous methods from literature will be analysed.

4.3.1 Random sample generation

This experiment serves to show the performance of a single simulation run. The PDF, moments frequency spectrum and autocorrelation function are analysed to ensure the correct behaviour is expressed. This is done to confirm that the CDF transformation method produces a PDF that has the correct shape, moments that agree with the theoretically calculated ones, a smooth monotonically decreasing frequency spectrum without large deviations from the theoretical shape, and an autocorrelation function that shows decorrelation to the extent that the result decays to zero. This experiment is a check before the more rigorous experiments are performed.

For this experiment a Log-Normal random number sequence has been generated, with the parameter σ set to 0.5, μ set to 0, and the bandwidth set to 0.025 Fs. The CDF transformation method discussed in this section was used to generate a correlated phase, and a correlated magnitude to the shape of a Log-Normal PDF. A total of one million correlated samples were generated.

Figure 157 shows the analytical PDF plotted against the histogram of the generated samples. There is a close fit between the analytical PDF and the measured samples. Figure 158 shows the measured moments obtained from the generated data, compared to the theoretical moments for the Log-Normal distribution. There is a close fit for the moments of lower order, and a small but negligible deviation is noticeable for the 6th order moment.

Performing the KL divergence test on the correlated Log-Normal random samples averaged over 30 experiments resulted in a value of 0.00010970. This value is 0.72% less than that of the correlated Log-Normal distribution generated in Section 3.6.2. This result is also less than that of both the Rayleigh distribution generated with the CDF transformation method, and the baseline method in Section 3.5.2.

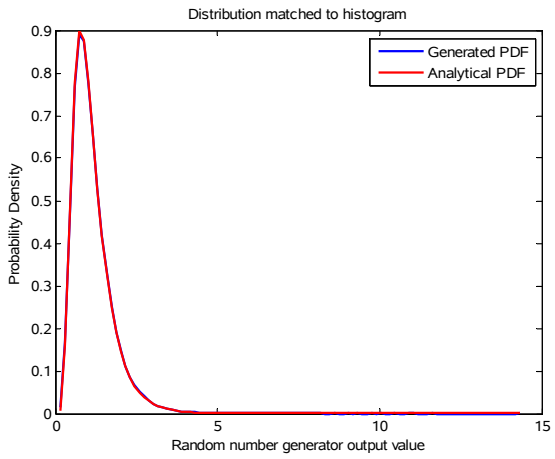

Figure 157: Normalised histogram of the time correlated Log-Normal random variable.

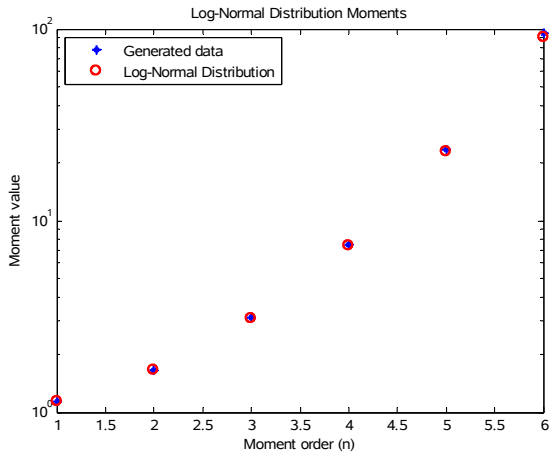

Figure 158: Measured moments from the generated correlated Log-Normal distribution.

Figure 159 shows the frequency spectrum of the correlated Log-Normal distributed samples generated with the CDF transformation method. The frequency spectrum is normalised to sampling frequency with the magnitudes taken relative to the largest frequency component. The frequency spectrum shows that there is correlation, but the measured bandwidth is larger than expected. The bandwidth was set to 0.025 Fs, but a value closer to twice that was measured. The frequency spectrum is smooth without a notch at 0.5 Fs. Figure 160 shows the autocorrelation function for the generated sequence. The sequence decorrelates completely at around 30 samples, which is much less than seen previously in Section 3.6.2.

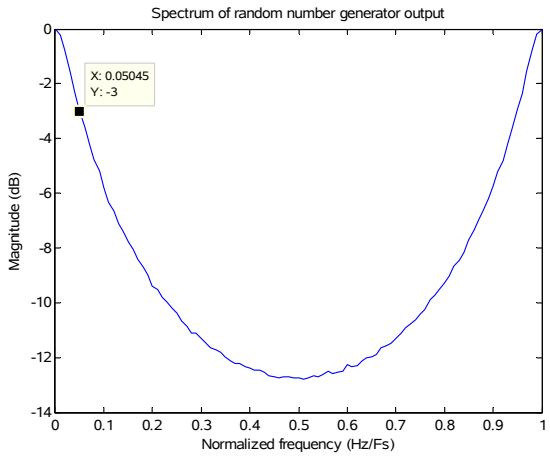

Figure 159: Frequency spectrum of the correlated Log-Normal random variable.

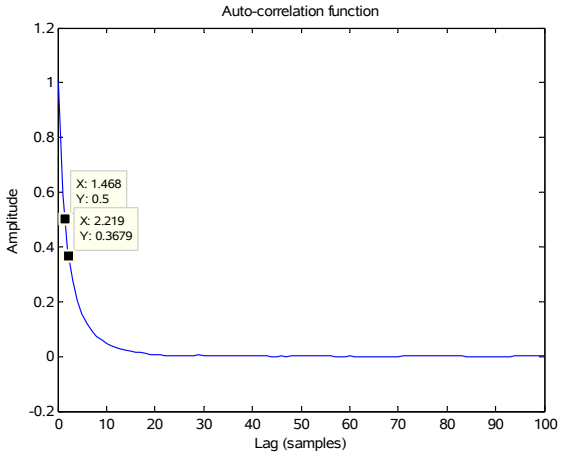

Figure 160: The autocorrelation function of the correlated random variable.

This experiment has shown that the correlated Log-Normal sequence generated with the CDF transformation method produces valid results. The Log-Normal PDF was visually accurate, and performed well when measured with the KL divergence method. The measured moments matched that of the theoretical ones. There are some inaccuracies observed in the bandwidth of the sequence where a larger value was observed with regards to the desired one, but the frequency spectrum shape was correct. Due to the duality between the time and frequency domain, the decorrelation times were less than expected. The experiments that follow will investigate this trend further.

4.3.2 Trend with varying correlation

This experiment investigates the trends in degradation in performance, by varying how strongly correlated the generated samples are. The desired bandwidth is stepped from 0.01 Fs to 0.25 Fs, with all other variables kept constant.

From the equivalent experiment in Section 3.6.3 the observed performance was favourable. The measured -3 dB bandwidth was very close to the desired -3 dB bandwidth, with a small error for small normalised bandwidth settings, which improved with an increase in the normalised bandwidth

setting. The quality of the PDF shape was also not acceptable with the χ^2 test rejecting the Log-Normal PDF just more than half of the time. The percentage error of the higher order moments were less than 6% across the entire range of bandwidths except for the degradation in performance for low bandwidth settings as was observed with the underlying Gaussian distribution.

Figure 161 shows the error between the set -3 dB bandwidths and the measured -3 dB bandwidths. Similarly to the Rayleigh distributed case implemented with the CDF method the measured normalised bandwidth has a large and non-constant error. As the normalised bandwidth setting increases the error becomes larger. Before the desired normalised bandwidth reaches 0.1 *Fs* the measured normalised bandwidth already crossed the 0.25 *Fs* mark. The results for the Log-Normal distributed samples are not much different to those of the Rayleigh distributed samples generated with the CDF transformation method.

Figure 162 shows the result of χ^2 tests, averaging over 100 experiments for each value of -3 dB bandwidth. It can be seen that the χ^2 test on average rejects the generated Log-Normal PDF. The performance for this method is very poor. Where previously the threshold of the χ^2 test and the average outcome was in the same range, there is now a factor 5 difference between the two. There is no trend in the data. Every single χ^2 test rejected the generated Log-Normal PDF. This performance is much worse than that of the equivalent experiment in Section 3.6.3.

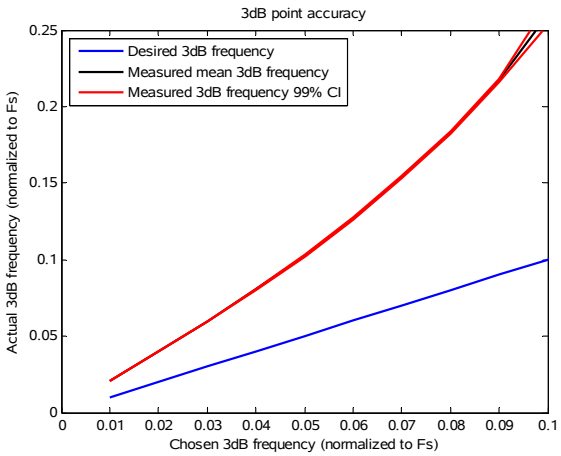

Figure 161: Measured -3 dB bandwidth of the correlated Log-Normal random number.

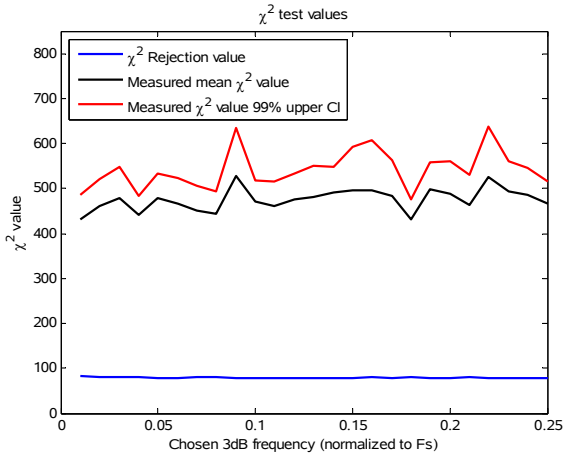

Figure 162: Plot of the outcome from a chi squared test.

Figure 163 shows the KL divergence value with varying degrees of correlation. There is a trend of worsening performance as the correlation in the samples are increased. There appears to be no increase in the variance of the results obtained except for the largest correlation at 0.01 *Fs* . The largest outlier worsens as the correlation is increased indicated by the red shot dashed line in the plot. Figure 164 shows the number of samples in the tail of the Log-Normal distribution, summed to

give a value out of 100 and Figure 165 shows the number of samples in the tail of the Log-Normal distribution, summed to give a value out of 10. There is no clear trend, and the number of samples in the tails appear to be constant irrespective of the normalised bandwidth setting. There is however an increase in variability of the result as seen by the variance shown on the plot by the blue lines, and the increase in maximum and minimum outliers shown on the plot by the red lines at the top and the bottom of the graph respectively. The KL divergence value for both the Rayleigh and the Log-Normal distributed samples generated with the CDF transformation method appears to be very similar. The only noticeable difference between the Rayleigh and the Log-Normal disturbed samples performed slightly worse for the strong correlation case at Fs equal to 0.01.

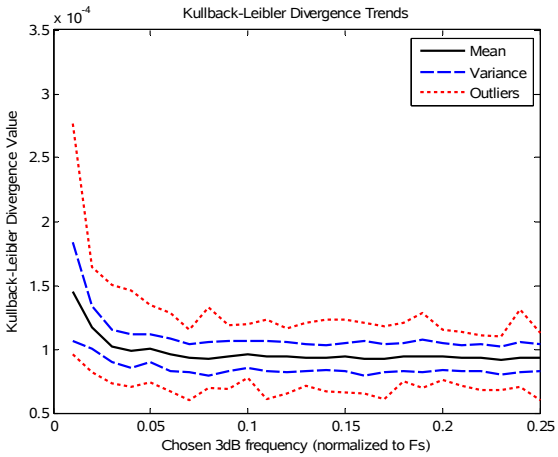

Figure 163: Kullback-Leibler divergence value.

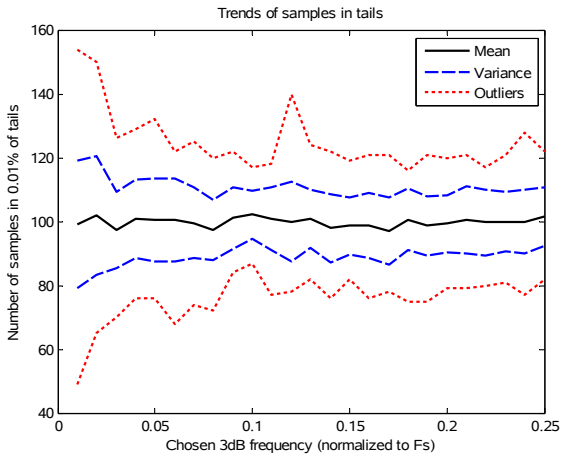

Figure 164: Number of samples in 0.01% of the tails.

Figure 166 shows the results from of the percentage error between the measured higher order moments and the calculated ones. All size the measured moments have the same trend with an error varying around zero percent irrespective of the bandwidth setting. The larger the order of the moment the larger the percentage error. The only visible trend is in the varying size of the errors which increases for smaller bandwidth settings, and smaller bandwidth settings. This increase at smaller bandwidth settings could be attributed to coincidence since none of the other tests showed this type of behaviour, and nothing in the CDF transformation method would cause this effect. The results here performed worse than that of the similar experiment in Section 3.6.3. These results are also approximately 10 times worse than the CDF transformation experiment performed with the Rayleigh distribution.

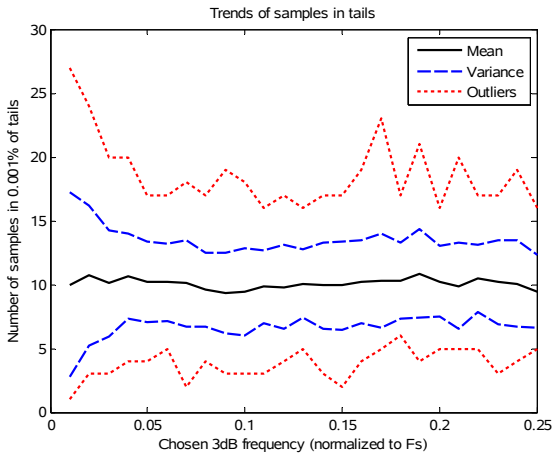

Figure 165: Number of samples in 0.001% of the tail.

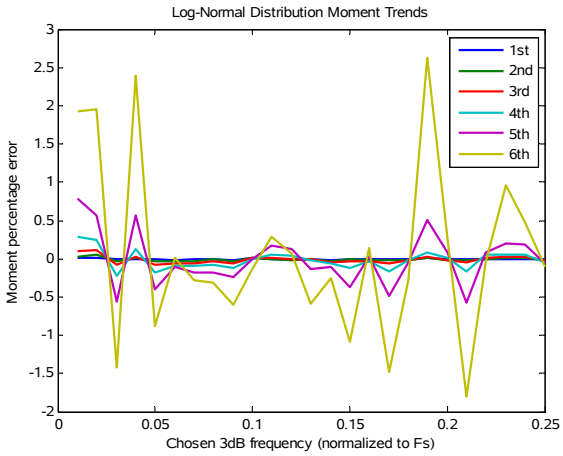

Figure 166: Percentage error of the measured moments from the generated correlated Weibull distribution.

This experiment has shown that the performance of the Log-Normal distributed samples generated with the CDF transformation method was similar to that of the Rayleigh distributed samples generated with the same method. The normalised bandwidth behaved in the same manner for both cases, and the higher order moments have shown the same trend, except for poorer performance for the Log-Normal than the Rayleigh. The KL divergence values were almost identical to that of the Rayleigh distributed samples. The χ^2 test however rejected the generated Log-Normal PDF outright without any indication of a trend in performance related to the normalised bandwidth.

4.3.3 Trend with varying Log-Normal scale parameter

This experiment investigates the trends in degradation in performance, by varying the Log-Normal scale parameter σ. For this experiment the value of the Log-Normal scale parameter was varied between 0.05 and 1 in steps of 0.05. The bandwidth of the signal was set to 0.025 Fs.

From the equivalent experiment in Section 3.6.4 a mixture between favourable and bad performance was observed. For small values of the Log-Normal scale parameter σ the performance was satisfactory, with the χ^2 test barely able to reject 10 out of 100 PDFs. As the Log-Normal scale parameter increased this performance remained relatively constant up to $\sigma=0.4$ where after a sharp drop in performance was observed. After $\sigma=0.55$ every PDF was rejected by the χ^2 test. The percentage error of the higher order moments showed the same trend, albeit with a more gradual degradation in performance.

The -3 dB accuracy is shown in Figure 167 for values of σ between 0.05 and 1. The trend observed is different from that observed for the Rayleigh distribution where the bandwidth was unaffected by the scale parameter. The result also does not show a similar trend to that of the equivalent

experiment in Section 3.6.4. For small values of the Log-Normal scale parameter σ the difference between the desired normalised bandwidth and the measured normalised bandwidth was less than that of larger values. The input bandwidth setting was kept constant at 0.025 Fs. At $\sigma=0.1$ the normalised bandwidth was measured as 0.045 Fs, and this value increases all the way up to $\sigma=1$ where the normalised bandwidth was measured as 0.064 Fs. There is a large error in the bandwidth, and more worryingly the error is non constant for different values of σ. Larger values of σ corresponds to a longer tail, and thus as the tail of the Log-Normal distribution becomes larger there is a larger decorrelation in the samples.

Figure 168 shows the result of χ^2 tests, averaged over 100 experiments for each value of σ. This result is similar to that obtained with the method from literature in Section 3.6.4 (Figure 94). It is clear that between $\sigma=0.4$ and $\sigma=0.5$ the performance of the distribution from the generated samples degrades drastically.

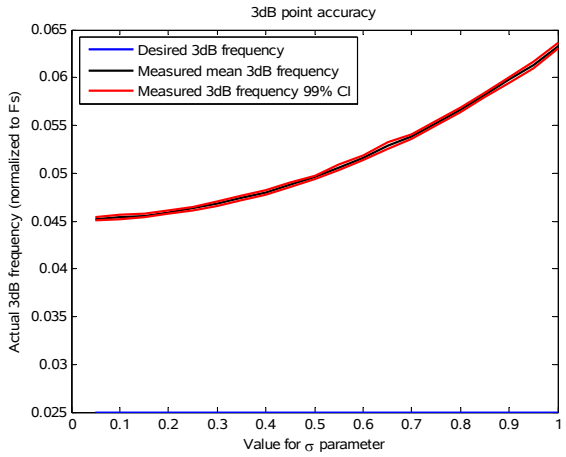

Figure 167: Measured -3 dB bandwidth of the correlated Log-Normal random number.

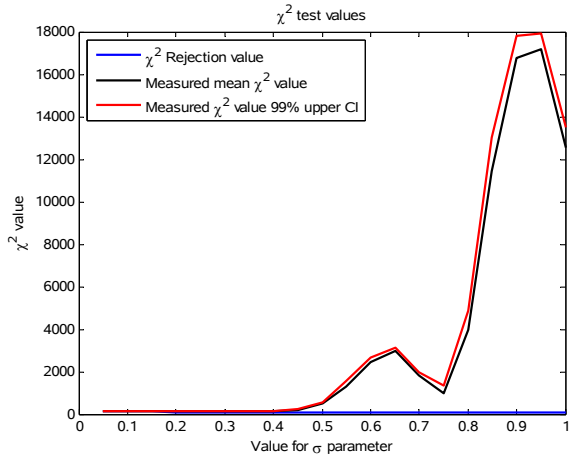

Figure 168: Plot of the outcome from a chi squared test.

Figure 169 is a zoomed in plot of Figure 168. From this figure it can be observed that the χ^2 test on average rejects the PDF of the generated Log-Normal random sequence for all values of sigma. This is poorer performance than that of the results in Section 3.6.4 where the performance only started to degrade for values of σ larger than 0.5. For that case the χ^2 test could not reject the PDF for values of σ less than 0.5. The sharp decrease in performance moved from 0.5 to 0.4 for the CDF transformation method, which is another indication of poorer performance with this method besides the overall poorer performance.

Figure 170 shows the number of χ^2 tests out of a 100 that were not rejected for each of the tested values of σ. The same trend is visible from this figure as from the previous two. There were at best more than 50 PDFs rejected for the best performing case at $\sigma=0.1$. From there on an increasing number of PDFs were rejected until ultimately at $\sigma=0.5$ all the PDFs were rejected. This performance is worse than that observed with the similar experiment in Section 3.6.4.

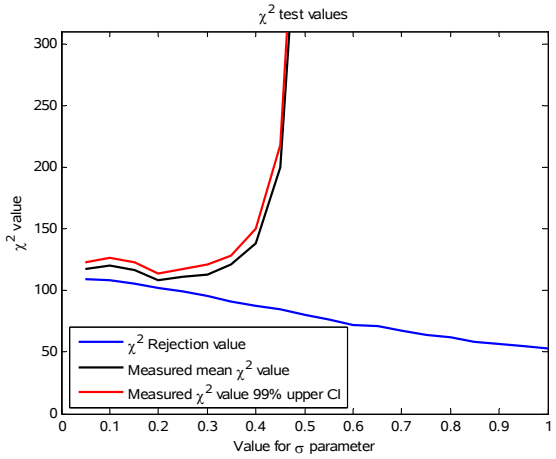

Figure 169: Plot zoomed in version of Figure 168 focusing on the trend of the chi-squared test threshold.

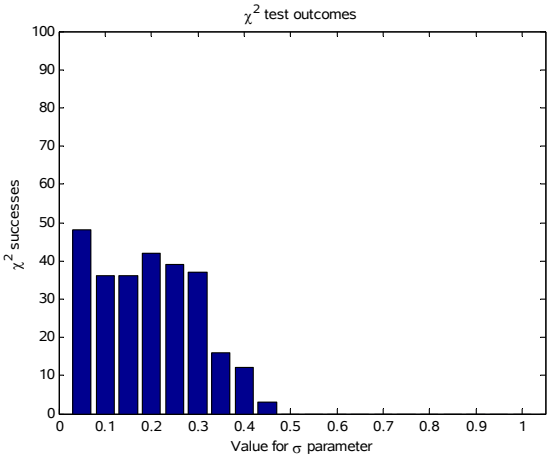

Figure 170: Number of the chi squared tests that were unable to reject the PDF.

Figure 171 shows the KL divergence value with varying values of the Log-Normal distribution scale parameter. There is a trend indicating improved performance for larger values of the scale parameter. The largest outlier, smallest outlier and variance show a slight increase in variability of the data as the scale parameter is increased. The same downward trend with increasing size of the scale parameter is seen with the Log-Normal distributed samples generated with the CDF transformation method than with the Log-Normal distributed samples generated with the method from literature in Section 3.6.4. The CDF transformation method however performs slightly worse than the method from literature.

Figure 172 and Figure 173 show the number of samples in the tail of the Log-Normal distribution. For Figure 172 0.01% of the tail is summed to give a value out of 100, and for Figure 173 0.001% of the tail is summed to give a value out of 10. There is no clear trend, and the number of samples in the tails appear to be constant irrespective of the Log-Normal scale parameter setting. The Log-normal distributed samples generated with the CDF transformation method had more variation in the results when compared to the method from literature in Section 3.6.4 for both the 100 sample and 10 sample cases.

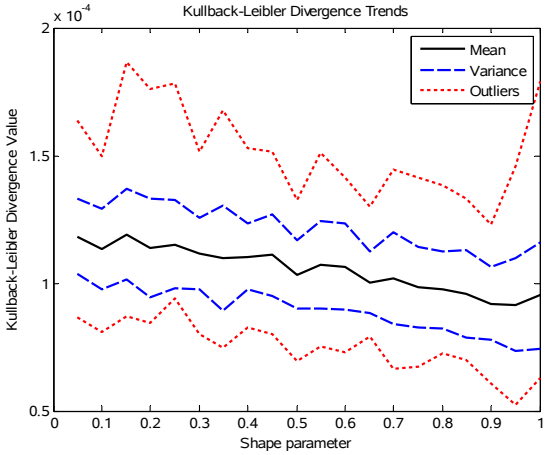

Figure 171: Kullback-Leibler divergence value.

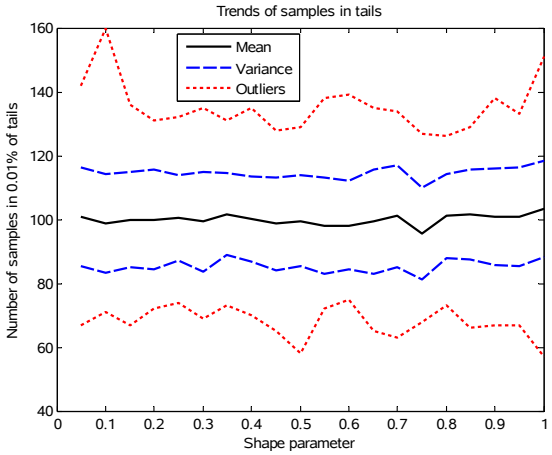

Figure 172: Number of samples in 0.01% of the tail.

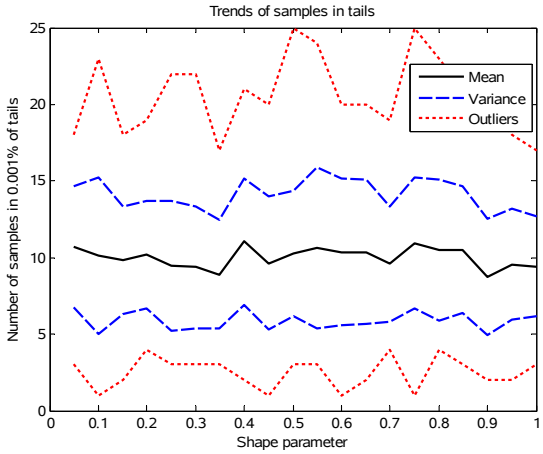

Figure 173: Number of samples in 0.001% of the tail.

The full scale of the percentage error for the higher order moments are shown in Figure 174 and a zoomed in version between -1% and 1% error is shown in Figure 175. The percentage error increases as σ increases. At approximately $\sigma = 0.6$ the 6th order moment percentage error declines sharply and continues to do so all the way to $\sigma = 1$. The percentage error in the moments does not show the same sharp switch between acceptable and poor performance that the χ^2 test has. The χ^2 test showed a clear change in performance at approximately $\sigma = 0.4$, where as the same degradation is observed at $\sigma = 0.6$ with the higher order moments. The percentage error of the moments do however indicate that the tail of the generated Log-Normal distribution is the source of the problems regarding the accuracy of the distribution, since the 5th and 6th order moments have shown the largest errors. The performance is worse for the CDF transformation method than the corresponding method from literature.

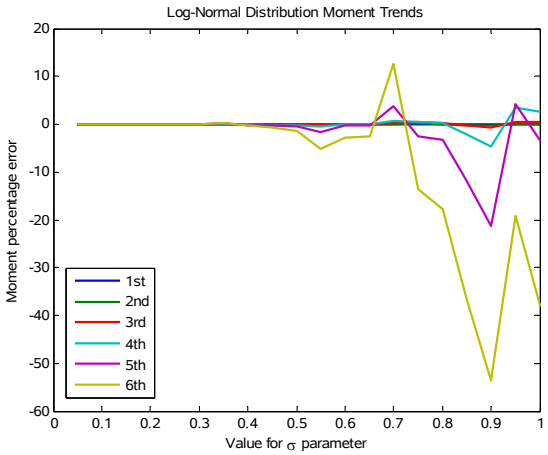

Figure 174: Percentage error of the measured moments from the generated correlated Log-Normal distribution.

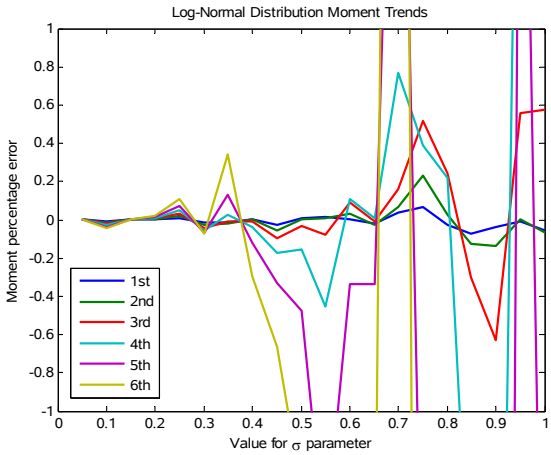

Figure 175: A zoomed version of the percentage error of the measured moments from the generated correlated Log-Normal distribution.

Similarly to the Rayleigh distribution generated with the CDF transformation method, the Log-Normal generated with the CDF transformation method showed poorer performance than that of the method in literature.

4.4 Weibull distributed cutter

The Weibull distribution was generated to serve as the magnitude of the clutter. An uncorrelated uniform variable is generated in the range between 0 and 1. The resulting random values are then passed through the inverted CDF of the Weibull distribution. The CDF of the Weibull distribution is shown in equation (56) and the inverse CDF in equation (57).

$$y = 1 - e^{-\left(\frac{x}{\lambda}\right)^{k}} \tag{56}$$

$$x = \lambda\left(-\ln\left(1-y\right)\right)^{\frac{1}{k}} \tag{57}$$

For brevity only the results that capture the overall performance, or highlights certain differences between this method and the previous methods from literature will be analysed.

4.4.1 Random sample generation

This experiment serves to show the performance of a single simulation run. The PDF, moments frequency spectrum and autocorrelation function analysed to ensure the correct behaviour is expressed. This is done to confirm that the CDF transformation method produces a PDF that has the correct shape, moments that agree with the theoretically calculated ones, a smooth monotonically decreasing frequency spectrum without large deviations from the theoretical shape, and an autocorrelation function that shows decorrelation to the extent that the result decays to zero. This experiment is a check before the more rigorous experiments are performed.

For this experiment a Weibull random number sequence has been generated, with the value k set to 1.5, λ set to 0, and the bandwidth set to 0.001 Fs. The CDF transformation method discussed in this section was used to generate a correlated phase, and a correlated magnitude to the shape of a Weibull PDF. A total of one million correlated samples were generated.

Figure 176 shows the analytical PDF plotted against the histogram of the generated samples. There is a close fit between the analytical PDF and the measured samples. Figure 177 shows the measured moments obtained from the generated data, compared to the theoretical moments for the Weibull distribution. There is a close fit for the moments of all tested orders.

The KL divergence value obtained when averaging over 30 experiments was 0.00011384 which is nearly identical to the corresponding experiment performed in Section 3.7.2. This value is 0.63% greater than that of the similar experiment performed with the method from literature. This is an indication that the CDF transformation method performs very well when only the KL divergence value is considered.

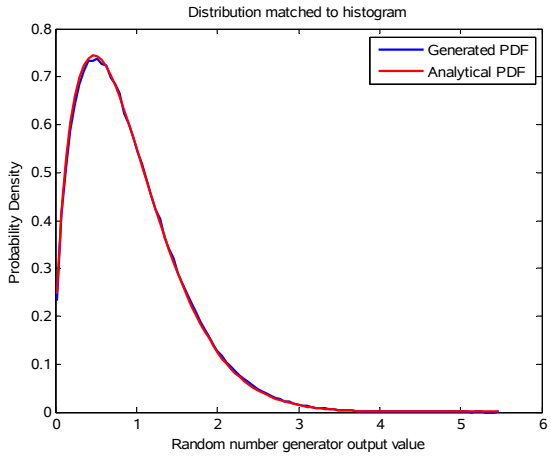

Figure 176: Normalised histogram of the time correlated Weibull random variable.

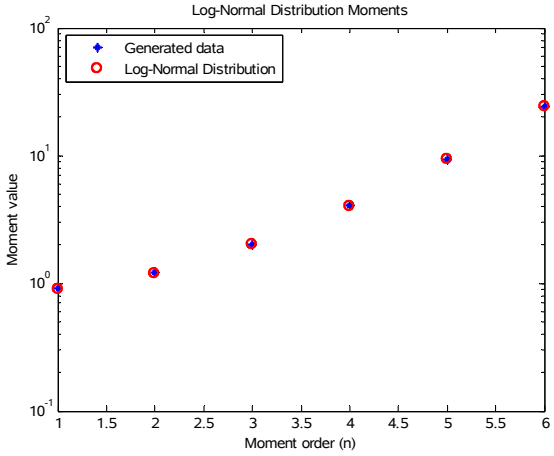

Figure 177: Measured moments from the generated correlated Weibull distribution.

Figure 178 shows the frequency spectrum of the correlated Weibull distributed samples generated with the CDF transformation method. The frequency spectrum is normalised to sampling frequency with the magnitudes taken relative to the largest frequency component. The frequency spectrum shows that there is correlation, but the measured bandwidth is larger than expected. The bandwidth was set to 0.01 Fs, but a value closer to twice that was measured. Figure 179 shows the autocorrelation function for the generated sequence. The correlation times were measured as 3.427 samples for $\tau_{1/2}$ and 5.295 samples for $\tau_{1/e}$. The sequence decorrelates completely at more than 30 samples, which is much less than seen previously.

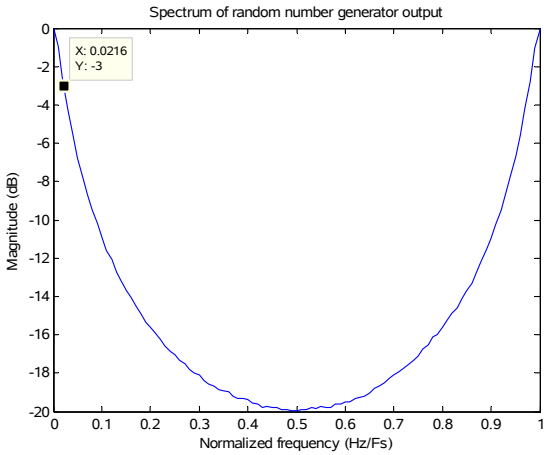

Figure 178: Frequency spectrum of the correlated Weibull random variable.

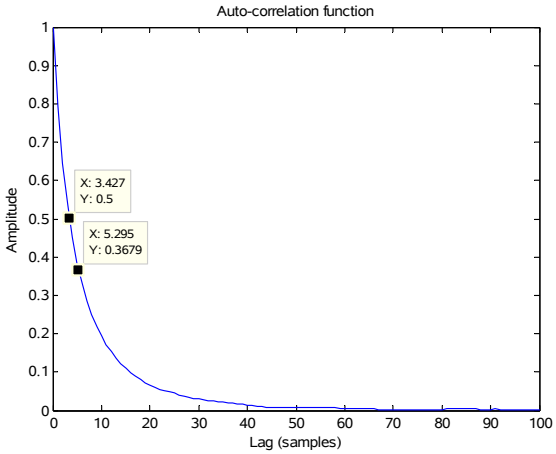

Figure 179: The autocorrelation function of the random variable.

This experiment has shown that the correlated Weibull sequence generated with the CDF transformation method produces valid results. The Weibull PDF was visually accurate, and the measured moments matched those of the theoretical ones. There are some inaccuracies observed in the bandwidth of the sequence, and thus due to the duality between the time and frequency domain, the decorrelation times were less than expected. The experiments that follow will investigate this trend further.

4.4.2 Trend with varying correlation

This experiment investigates the trends in degradation in performance, by varying how strongly correlated the generated samples are. The desired bandwidth is stepped from 0.01 Fs to 0.25 Fs, with all other variables kept constant.

From the equivalent experiment in Section 3.7.3 the performance observed was not satisfactory. Although the measured normalised bandwidth very closely matched that of the desired normalised bandwidth, the χ^2 test rejected all generated of the Weibull PDFs, and the average χ^2 outcome was much larger than the threshold. The percentage error of the higher order moments were best for small bandwidth settings and increased quickly for larger bandwidth settings. The reason for the poor performance was failure of the method to generate a PDF that resembles that of a Weibull distribution, and increasingly so for larger values of bandwidth.

Figure 180 shows the difference between the set -3 dB bandwidths and the actual measured -3 dB bandwidths. Similarly to both the Rayleigh and Log-Normal distributed cases implemented with the CDF transformation method the measured normalised bandwidth has a large and non-constant error. As the normalised bandwidth setting increases the error becomes larger. Before the desired

normalised bandwidth reaches 0.1 Fs the measured normalised bandwidth had exceeded 0.25 Fs. The results for the Weibull distributed samples are not much different to those of the Rayleigh and Log-Normal distributed samples generated with the CDF transformation method.

Figure 181 shows the result of χ^2 tests, averaging over 100 experiments for each value of -3 dB bandwidth. It can be seen that the χ^2 test on average rejects the generated Weibull PDF. The performance for this method is different to that observed in Section 3.7.3 and similar to the performance trend of the Rayleigh distributed samples generated with the CDF transformation method. There is a decrease in performance when the desired bandwidth is set lower than 0.05 Fs, but remains relatively constant for all bandwidth settings larger than this value.

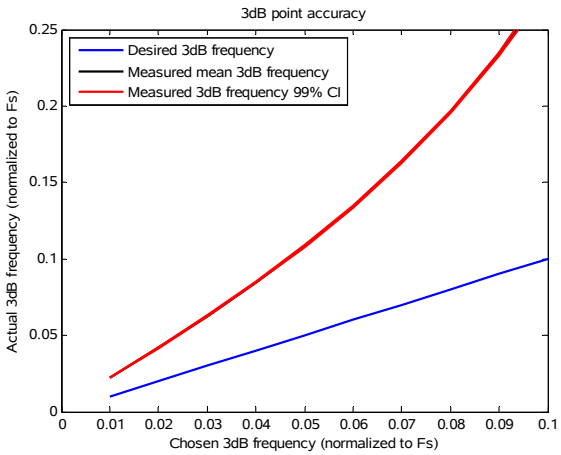

Figure 180: Measured -3 dB bandwidth of the correlated Weibull random number.

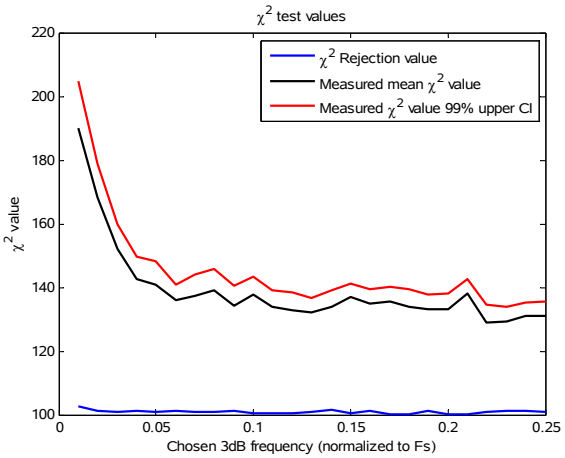

Figure 181: Plot of the outcome from a chi squared test.

Figure 182 shows the number of χ^2 tests out of a 100 that were not rejected across the normalised bandwidth range tested. The average number of tests not rejected were never more than 10. There is no clear trend visible from this result. This confirms the poor result seen in Figure 181.

Figure 183 shows the KL divergence value with varying degrees of correlation. There is a trend of worsening performance as the correlation in the samples are increased. There appears to be no increase in the variance of the results obtained. The largest outlier worsens as the correlation is increased indicated by the red shot dashed line in the plot. The average result is roughly the same to that of the similar experiment performed on the Weibull distributed sample generation method from literate from Section 3.7.3.

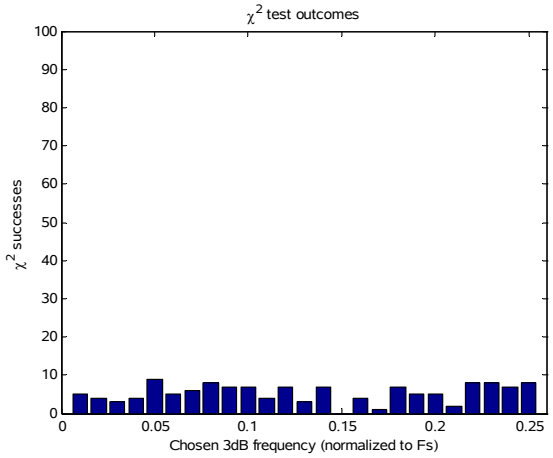

Figure 182: Number of chi squared tests that were unable to reject the PDF.

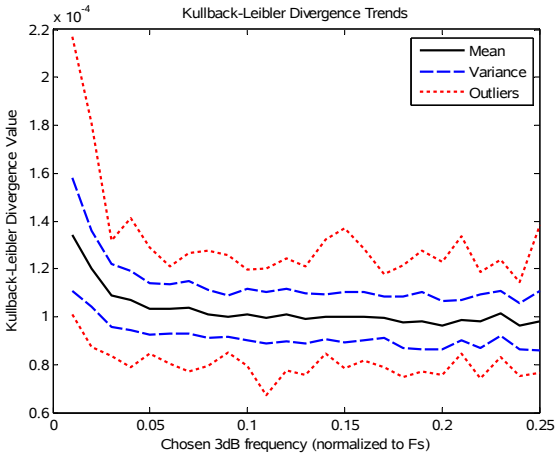

Figure 183: Kullback-Leibler divergence value.

Figure 184 shows the number of samples in the tail of the Weibull distribution, summed to give a value out of 100. There is no clear trend, and the number of samples in the tails appear to be constant irrespective of the normalised bandwidth setting. There is however an increase in variability of the result as seen by the variance shown on the plot by the blue lines, and the increase in maximum and minimum outliers shown on the plot by the red lines at the top and the bottom of the graph respectively.

Figure 185 shows the results from of the percentage error between the measured higher order moments and the calculated ones. All size the measured moments have the same trend with an error varying around zero percent irrespective of the bandwidth setting. The larger the order of the moment the larger the percentage error. This result as obtained with the CDF transformation method compares favourably to the similar result obtained with the method from literature. This result is also better than the results obtained with from the Log-Normal samples generated with the CDF transformation method.

This experiment has shown that the performance of the Weibull distributed samples generated with the CDF transformation method was similar to that of the Rayleigh and Log-Normal distributed samples generated with the same method. The normalised bandwidth behaved the same manner for all three cases, and the higher order moments has shown the same trend, except for poorer performance for the Weibull than the Rayleigh. The χ^2 test however rejected the generated Weibull PDF outright and the same trend in performance related to the normalised bandwidth was observed here as was observed for the Rayleigh distribution.

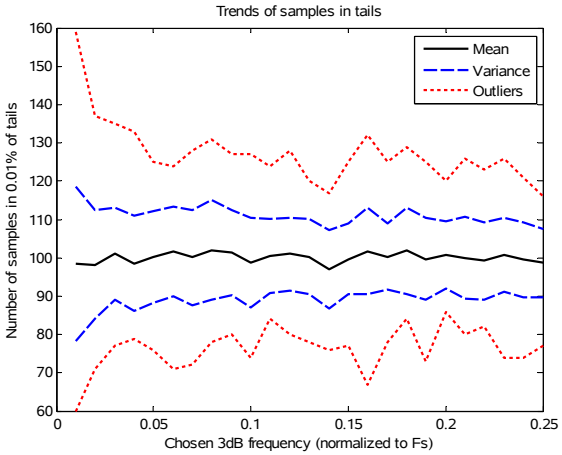

Figure 184: Number of samples in 0.01% of the tail.

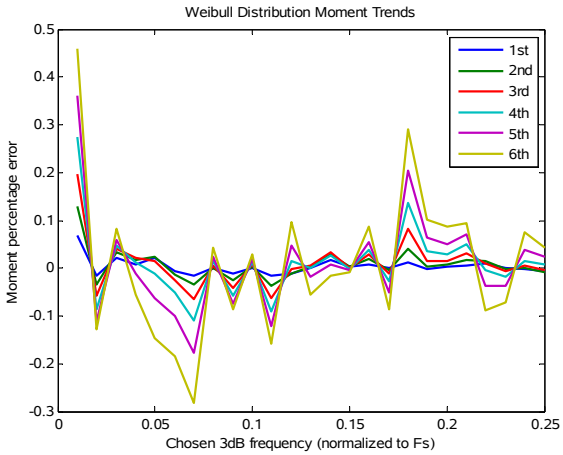

Figure 185: Percentage error of the measured moments from the generated correlated Weibull distribution.

4.4.3 Trend with varying Weibull shape parameter

This experiment investigates the trends in degradation in performance, by varying the Weibull shape parameter k. For this experiment the value of the Weibull scale parameter was varied between 1 and 5. The bandwidth of the signal was set to 0.01 Fs.

From the equivalent experiment in Section 3.7.4 the performance was relatively poor. For small values of the Weibull shape parameter k there was a larger decorrelation between samples and thus a larger normalised bandwidth measured at the smaller values than measured for the larger values. The χ^2 test on average rejected the Weibull PDF for all values of k showing a sharp peak where k is close to 1.2. Besides the peak the performance of the χ^2 test did not indicate any trend. Overall more than 90 out of the 100 χ^2 tests rejected the Weibull PDF. The Higher order moments showed worse performance for small values of k than for large values. The degradation in performance was gradual with no clear trend.

Figure 186 shows the result of the -3 dB bandwidth of the generated Weibull sequence as a function of the k parameter. The trend observed is different from that observed for the Rayleigh distribution where the bandwidth was unaffected by the scale parameter. The result shows a similar trend to that of the equivalent experiment in Section 3.7.4. The only difference is that the measured normalised bandwidth is much larger for the CDF transformation method than for the method from the literature. This result has been analysed in the previous experiment. For small values of the Weibull shape parameter k the difference between the desired normalised bandwidth and the measured normalised bandwidth was greater than that of the larger values. The input bandwidth setting was kept constant at 0.01 Fs. At $k=1$ the normalised bandwidth was measured as 0.0245 Fs, and

this value decreases all the way to $k=5$ where the normalised bandwidth was measured as 0.0195 Fs. There is a large error in the bandwidth, and more worryingly the error is not constant for different values of k. Smaller values of k corresponds to a longer tail, and thus as the tail of the Weibull distribution becomes larger there is a larger decorrelation in the samples. This same result was observed for the Log-Normal distribution.

Figure 187 shows the result from the χ^2 tests. On average not a single value for k had a successful average χ^2 test outcome. The Weibull random sequence generator PDF performed worse at $k=1.2$, only improving slightly at $k=1.4$. For $k=1$ and k larger than 1.6 the performance was better, and roughly the same, but sill not below the required threshold. This result is strikingly similar to that of the corresponding experiment in Section 3.7.4. This could indicate that the shape of the PDF has a large influence on the outcome of the χ^2 test with respect to the threshold value obtain for the distribution.

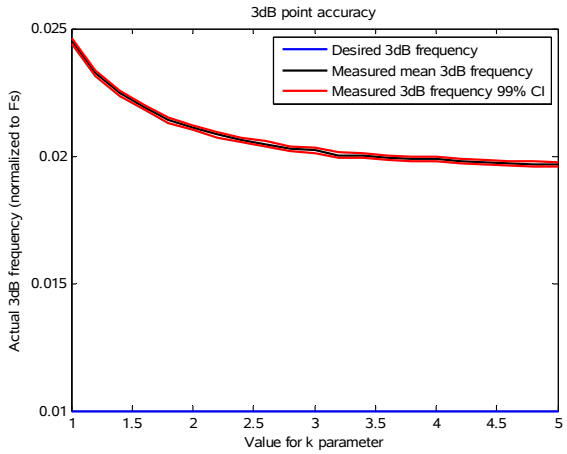

Figure 186: Measured -3 dB bandwidth of the correlated Weibull random number.

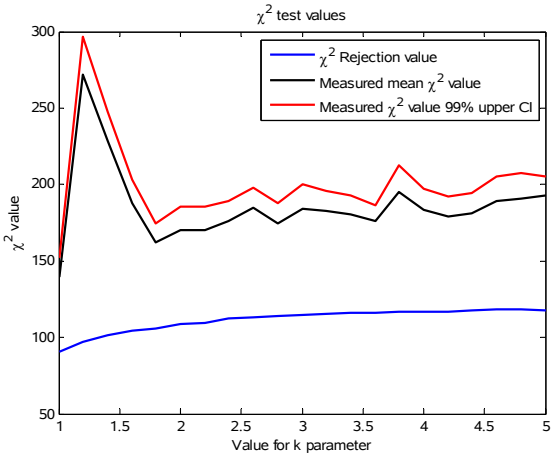

Figure 187: Plot of the outcome from a chi squared test.

Figure 188 shows the histogram of the result from the χ^2 tests. Out of 100 the more than 90 Weibull PDFs were rejected by the χ^2 test for most of the values of k. The performance is slightly better than that of the equivalent experiment in Section 3.7.4 but still poor.

Figure 189 shows the KL divergence value with varying values of the Weibull distribution scale parameter. There is no trend in the performance as the shape parameter k is varied. The largest outlier, smallest outlier and variance show no observable trend in the KL divergence value.

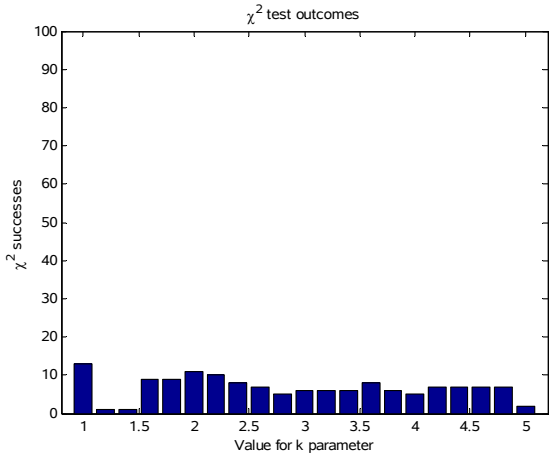

Figure 188: Number of chi squared tests that were unable to reject the PDF.

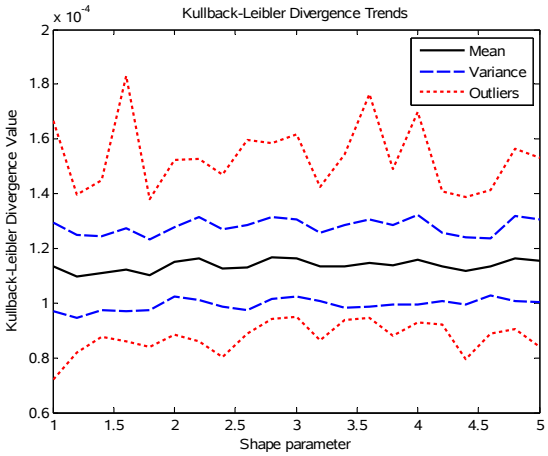

Figure 189: Kullback-Leibler divergence value.

Figure 190 shows the number of samples in the tail of the Weibull distribution, summed to give a value out of 100. There is again no trend, and the number of samples in the tails appear to be constant irrespective scale parameter setting. Unlike for the previous experiment where the correlation was varied, there is no change in variability of the result as seen by the variance shown on the plot by the blue lines.

The percentage error of the higher order moments is shown in Figure 191. The higher order moments show a trend of improvement with an increase in the shape parameter k. The improvement can mostly be noticed in the variation of the result. The 6th order moment always showed the largest error, except for the case where k is equal to 1, where all the moment errors were positive, and the 6th order moment negative. This could be due to the fact that the Weibull distribution loses samples in its tail.

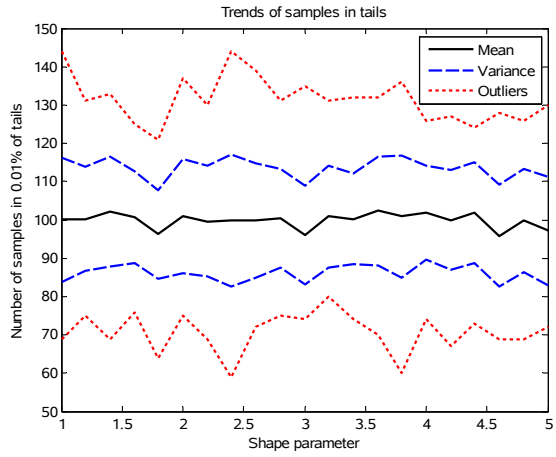

Figure 190: Number of samples in 0.01% of the tail.

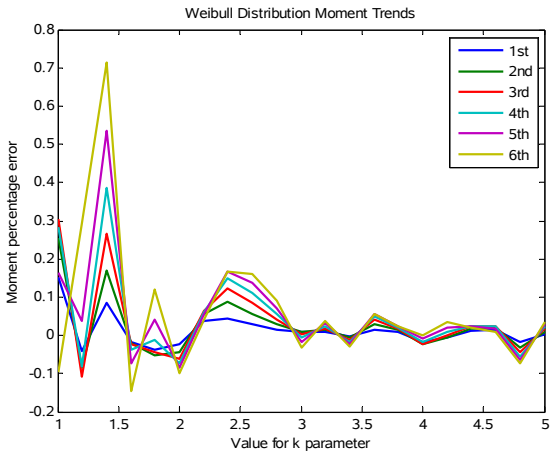

Figure 191: Percentage error of the measured moments from the generated correlated Weibull distribution.

The Weibull distributed samples generated with the CDF transformation method showed similar performance for the χ^2 test, and improved performance for the higher order moments. The normalised bandwidth was worse but similar to the observed with the Rayleigh and Log-Normal distributed samples.

4.5 Conclusion

The three distributions commonly used for ground clutter modelling were investigated in Section 3 and were again analysed here for the new CDF transformation method for clutter simulation. The three distributions were the Rayleigh, Log-Normal and Weibull distributions. The CDF transformation method uses a combination of CDF and ICDF transformations to obtain a correlated uniform phase and a correlated magnitude with the desired PDF. The CDF transformation method starts off with two correlated Gaussian distributed sequences, both with the same bandwidth setting. The two sequences with Gaussian distributed samples were then transformed to two sequences with uniformly distributed samples by feeding the values through the Gaussian CDF. One of the correlated uniformly distributed sequences is then used for the phase after shifting and scaling, while the other is transformed to the required magnitude distribution through the corresponding ICDF of that distribution. This represents a signal in polar coordinates. The last step is the conversion to Cartesian co-ordinates which produces the I and Q channel signals.

The method used to generate the Gaussian distributions was the same method as discussed in Section 3.2. The phase made use of a uniform distribution transformation, and it showed performance that matched closely the performance of the Gaussian distribution used as input. The distribution was flat, and the frequency spectrum and autocorrelation function showed near identical results to that of the Gaussian distribution. Measuring the performance for varying degrees of correlation showed that for narrow bandwidth settings the error for the Gaussian and uniformly distributed signals were the same. As the bandwidth increases the error for the Gaussian distribution becomes smaller but the error for the uniform distribution increases slightly. The error is still very small for either case. There is thus a negligible difference between the Gaussian and uniformly distributed correlated sequences, which is a positive result. This behaviour was not observed when the uniformly distributed samples were further transformed to other distributions.

The transformation of the uniform distribution to the Rayleigh, Log-Normal and Weibull distributions added decorrelation in the generated sequences. The experiments showed large errors between the desired bandwidth and the actual bandwidth. Settings between 0.01 Fs and 0.1 Fs resulted in bandwidths between 0.02 Fs and 0.25 Fs. The error was not a straight line, indicating

that addition or multiplication by a constant factor alone would not be able to completely remove it. The solution is thus more involved and requires more attention. All three distributions exhibited the same trend, but these did not match exactly and thus the decorrelation was found to be dependent on the distribution of choice. This was confirmed when the scale parameter of the Log-Normal and shape parameter of the Weibull distribution were varied (the scale parameter of the Log-Normal also changes its shape). As the tails of a distribution become larger relative to the body of the distribution, the decorrelation increases. It makes sense then, when varying the scale parameter of the Rayleigh distribution, no effect on the decorrelation was measured, because the underlying shape is not altered, it is merely scaled. There was however a one-to-one relationship with the CDF transformation method, where a single input bandwidth produced a single output bandwidth, which means it could potentially be corrected for with curve fitting.

The χ^2 test did not reject the Rayleigh PDF generated with the CDF transformation method for most bandwidth settings. The performance was relatively flat and started to degrade for bandwidth settings less than 0.1 Fs. The generated Log-Normal distributions were rejected by the χ^2 test for every experiment performed. The generated Weibull distributions were rejected by the χ^2 test in more than 90% of the experiments. The Log-Normal did not show any trend in the χ^2 test results by varying the bandwidth, and the Weibull distribution showed the same trend as the Rayleigh distribution, with degradation in performance for narrower bandwidth settings. The KL divergence value showed very similar results for the CDF transformation method, when compared to the methods analysed from literature in Section 3. This is an indication that the KL divergence value, although useful for the comparison of a method with itself with varying parameters, it is not as useful when comparing distributions generated with different methods. The χ^2 test provided much more useful results than the KL divergence value.

The percentage error for the higher order moments showed the same trend for all three distributions, with a larger variation in error for narrow bandwidth settings and a relatively unchanged error from 0.1 Fs upward. The Rayleigh distribution performed the best with its largest percentage error of 0.15% for the 6th order moment. The Weibull distribution performed very similarly but slightly worse with its largest percentage error around 0.4% for the 6th order moment. The Log-Normal distribution performed very badly compared to the other two. The percentage error of the 6th order moment averaged around 2.5% at its worst.

Keeping the bandwidth constant to measure the effect of the shape parameter of a distribution on the accuracy of its PDF had no unexpected results. There was no variation in the higher order moments by varying the Rayleigh scale parameter. The 6th order moment remained below a 0.25% error. The Log-Normal distribution showed poorer performance for larger values of the Log-Normal scale parameter. The performance was relatively constant for values less than 0.4 of the scale parameter, but after this point the performance degraded drastically and continued to do so all the way to the largest value tested for the scale parameter of 1. The performance of the percentage error for the higher order moments was relatively positive when the Weibull shape parameter was varied. The worst performance observed was just above a 0.7% error for the shape parameter value of 1.4. The percentage error remained below 0.2% for values of the shape parameter between 2 and 5. Both the Log-Normal and Weibull distributions showed poorer performance when the distributions had longer tails.

The χ^2 test did not show any trend when the scale parameter of the Rayleigh distribution was varied, but the tests did reject more than 50% of the generated distributions. The Log-Normal generated samples showed the same curve where the χ^2 test showed better performance for small values of the scale parameter, and poorer performance for larger values of the scale parameter. At best the χ^2 test rejected 50% of the generated Log-Normal distributed samples, and after a value of 0.4 for the scale parameter, all the distributions generated with the CDF transformation method were rejected. For the generated Weibull distribution the χ^2 test showed the same performance except for an anomaly at values of 1.2 for its shape parameter. These results match closely those from the previous section, although the performance from the CDF transformation method was worse. This indicates that most of the performance variation observed with the χ^2 test might be due to the actual shape of the distribution at this setting, and a difficulty for the χ^2 test to measure these shapes, rather than poor performance for those specific generated distribution shapes.

The performance was unfortunate though, since the poorest performance for all the generated distribution shapes was observed for narrow bandwidths, and to correct for the error in the bandwidth, the input setting will have to be adjusted further into the region of smaller bandwidths. This is not as much of a problem with the methods from literature as seen in Section 3, because a negligible bandwidth error exists for those cases.

The distributions however, did not suffer any of the drawbacks that were noticed from the methods implemented in Section 3. The Log-Normal and Weibull distribution were much easier to generate with the CDF transformation method than the methods from literature. The shapes of the

distributions held up fairly well, even though the χ^2 tests rejected most of them. This can be attributed to an overly high confidence level chosen for these experiments. Since clutter is only modelled by these distributions, a very accurate distribution cannot be a strict requirement.

The results from the CDF transformation method were thus mixed. The problems observed for the methods from literature were not evident with this approach, but predictability and control over the bandwidths of the final distribution have been sacrificed. The performance of the shape of the PDF was also worse than with the methods from literature. The CDF transformation method thus traded better accuracy for less complexity. There is still further research to be performed to enhance and optimise the CDF transformation approach for clutter simulation on a FPGA based platform.

5 CONCLUSION

The literature survey in Section 2 highlighted 8 different properties of radar clutter that have to be taken into account when clutter is simulated. The 8 different properties of clutter were:

1. RCS,
2. number of discrete scatterers,
3. spatial extent,
4. velocity extent,
5. wavelength,
6. amplitude distribution,
7. spatial correlation, and
8. polarization.

The clutter RCS scales the mean magnitude of the return and is derived from the geometry of the simulated scenario (angle, distance and reflectivity of the scatterer), while the number of discrete scatters is captured within the amplitude distribution of the clutter. When the behaviour of clutter is analysed over an entire range line, or for a scanning radar, the spatial extent of the clutter is important, but it does not play a role within the simulation of a single patch of clutter, which was the focus of this project. The velocity extent of clutter is mostly affected by the movement of a radar, but if the radar and its antenna are stationary, the movement of the clutter itself (for example wind blown trees) is the dominant effect. The correlation properties of clutter are strongly related to the velocity (spectral) extent of the clutter. The wavelength of radar clutter mostly influences the magnitude of the clutter return, and the behaviour of clutter, which is also captured by the amplitude statistics of the simulation.

To evaluate the simulated clutter, various tests were performed. To evaluate the amplitude statistics of the simulated clutter, the generated PDFs were measured, and plotted against their analytical equivalents. The error measured between the two was also plotted. In addition to those two measures the χ^2 test was used in combination with confidence intervals to evaluate the accuracy of the generated distributions. The KL divergence method was used as a statistically valid distance measure between the analytical and simulated distributions. It was however only useful when comparing the accuracy of a distribution with itself, and at times showed results that contradicted the results from the χ^2 test (for example when the tails of the Log-Normal distribution were much larger than the body of the distribution). An FFT was performed to evaluate the frequency spectrum of the simulated data, and the autocorrelation function was used to evaluate the correlation of the

simulated clutter. The higher order moments of distributions are often used to evaluate the accuracy of the tail regions of PDFs. For radar clutter the tail of the distribution is very important, and the higher order moments (up to the 6th order) have been included as a measure of the accuracy of the simulated clutter distributions.

Three distributions were identified as the most commonly used for ground clutter simulation. These distributions were the Rayleigh, Log-Normal and Weibull distributions. In Section 3 these distributions were simulated with methods found in literature that could potentially be implemented on to an FPGA with little modification. The MNLT methods were chosen because they produce samples sequentially, rather than in batches as is the case with the SIRV type methods. The MNLT methods use a correlated Gaussian distribution, and transform it to the required distribution. The Gaussian distribution thus formed the basis of all the tested distributions and its performance was evaluated first.

The correlated Gaussian distribution performed well. The shape of the distribution was measured to be accurate and the measured correlation was very close the chosen values. For narrow bandwidth settings the Gaussian distribution did show signs of a degradation in performance. The PDF accuracy declined and the measured bandwidth were larger than the chosen value. These discrepancies however were negligibly small. The shape of the frequency spectrum and the shape of the autocorrelation function did not show any unexpected behaviour.

The MNLT method was not used for the generation of the Rayleigh distribution, instead the correlated Gaussian distributed samples were used as is for the I and Q channel. The magnitude of the two Gaussian distributions produced the correlated Rayleigh distribution. The resultant Rayleigh distribution performed better than the two underlying Gaussian distributions. The frequency spectrum showed the expected shape. The autocorrelation function decayed all the way down to zero as it should. The phase of the generated samples was within the correct range and uniformly distributed. The PDF was more accurate and the percentage error in the higher order moments was very small. The error for small bandwidth values was still present, and the performance of the PDF still degraded for the small bandwidth settings. The scale parameter of the Rayleigh distribution had no detrimental effects on the performance of the distribution.

The Log-Normal distributed clutter was generated with an MNLT method. The MNLT method used two correlated but independent Gaussian sequences as input, and transformed them to the required Log-Normal distribution. The resultant Log-Normal distribution performed well. The method from [38] was implemented because it was an improvement on the method of [39] which had a Gaussian phase distribution instead of the required uniform phase distribution. The frequency spectrum

showed the expected shape. The autocorrelation function decayed all the way down to zero as it should. The phase of the generated samples was within the correct range and uniformly distributed. The Log-Normal PDF was found to be less accurate than the Rayleigh distributed samples, and the percentage error in the higher order moments was much larger. The error for narrow bandwidth values was still present, and the performance of the PDF still degraded for the small bandwidth settings. The scale parameter for the Log-Normal distribution affects the shape of the distribution, and had a large effect on the performance of the distribution. For small values of the scale parameter the performance of the Log-Normal distributed samples passed the PDF shape tests, but for larger values it failed drastically. This failure was attributed to the size of the tails of the distribution caused by the scale parameter. This method to generate Log-Normal samples was found to be more complex than the method chosen for Weibull distributed clutter generation, and would be difficult to implement on a hardware platform because of the many mathematically complex operations required.

For the Weibull distribution the method from [42] was implemented. This method did not suffer as much from the computational complexity drawbacks as much as the one in [38]. The frequency spectrum was clean and showed the expected shape. The autocorrelation function decayed all the way down to zero as it should. The phase of the generated samples was within the correct range and uniformly distributed. Variation in the correlation for the distribution had the same results as was observed for the Gaussian and Rayleigh distribution, with little error for wider bandwidth settings and a small but negligible error when the bandwidth settings became narrow. The χ^2 error and the higher order moments which tested the PDF shape, showed poor performance. The χ^2 test rejected almost every generated PDF. The higher order moments did not show large errors however. These two tests also showed a degradation in performance as the bandwidth was decreased.

Because of the complexity of the Log-Normal distribution generation algorithm, and the poor PDF performance of the Weibull distribution generation algorithm, neither of the methods chosen would be usable on an FPGA platform. This prompted the question as to weather there was a different method that could be designed, that from the start would map easily onto an FPGA. Why is clutter not simulated by creating a uniform phase, and the required magnitude as the radar would observe it? A new method using the CDF and ICDF of a distribution was designed and tested to answer this question in Section 4.

The CDF transformation method used a combination of CDF and ICDF transformations to obtain a correlated uniform phase and a correlated magnitude with the desired PDF. The CDF transformation method starts off with two correlated Gaussian distributed sequences, both with the same bandwidth

setting. The two sequences of Gaussian distributed samples were then transformed to two sequences with uniform distributed samples by passing the values through the Gaussian CDF. One of the correlated uniformly distributed sequences is then used for the phase after scaling, while the other is transformed to the required magnitude distribution by passing the values through the corresponding ICDF of that distribution. This represents a signal in polar coordinates. The last step is a conversion to Cartesian co-ordinates which produces the I and Q channel signals.

The CDF transformation method resolved some of the problems observed with the PDF of the generated samples for the Weibull distribution. The frequency spectrum was clean exhibiting the correct shape, for all three the tested distributions. The generated PDFs did not degrade in performance as was observed with the methods [42] and is much more compatible with FPGA implementation than the method form [38]. These improvements did come with some drawbacks though. The bandwidth of the correlated samples was no longer easily controllable because the input bandwidth set-point for the Gaussian distributions did not match the measured bandwidth of the generated samples with the transformed distributions. The trends in performance of the PDF shape for the CDF transformation method were similar to the methods from Section 3 but the overall performance measured was worse. The performance for all the distributions however showed the same trends, with worse performance for longer tails, and narrower bandwidth settings. This consistency in the results across different distributions is a favourable property to have for a method destined for FPGA implementation.

This new CDF transformation method thus had mixed results. Although the complexity problems observed using the methods in Section 3 were resolved additional problems were introduced with regards to the accuracy of the results. There is still room for further investigation into this method, to determine whether these problems can be solved. The transformation to the uniform distribution has little effect on the correlation, but the additional step from the uniform distribution to the desired magnitude distribution has severe loss in sample to sample correlation. How would the -3 dB bandwidth of the spectrum be affected if the correlation of the magnitude is adjusted separately from the phase, to match that of the required correlation? Would this eliminate the problem altogether? There is also the question regarding the effect of the size of the correlation of the magnitude. Some experiments during the development of the simulations showed to a certain extent, that correlation in the magnitude is not always required, which could mean that the correlation of the magnitude is not responsible for the loss in the correlation observed for this

method, but rather the shape of the magnitude. Exactly what the conditions for this are, has not been investigated. Additionally, it might be possible to control the shape of the spectrum and the PDF with different bandwidths for the magnitude and phase.

Further research is required to determine if the extent of the decorrelation can be determined beforehand. It was observed that the larger the tails of the distribution, the larger the decorrelation. If the value of the higher order moments, or position of a certain carefully chosen combination of percentiles of a PDF can be used to predict the expected decorrelation, one could possibly correct for it given any distribution, allowing the user more freedom in the choice of clutter statistics. Pre-calculation of the decorrelation is also an option but this can be time consuming due to the large number of experiments that has to be completed to reduce the inaccuracy of finite random sample sets sufficiently. The pre-calculation step is a brute force approach, and if a mathematically derived relationship for certain distributions can be found it would be of great value.

The attractive aspect of the CDF transformation method is that a single input passed through the correct transformation function produces the required output. This transformation function is a small piece of the algorithm that has to be changed when different distributions are to be simulated. The CDF transformation method can be implemented numerically which has certain advantages. The benefit of this is that any conceivable distribution can be simulated with this method. If a user would like to simulate non-standard clutter amplitude statistics observed during radar operation, that cannot be captured with an analytical function, it would be entirely possible. The fact that there already exists an implementable solution to solve the bandwidth problem, even if at this stage it is still very crude, means that the CDF transformation method is useful for FPGA implementation. A lookup table can correct the bandwidth problems when the required distribution is loaded and configured, and the behaviour of the frequency spectrum, autocorrelation function and PDF shape is within acceptable ranges.

To measure the accuracy of a PDF is a standard problem in statistics. The χ^2 test is just one of the methods used to measure the accuracy of a PDF. The main problem with the χ^2 test that was observed for the experiments performed here was that the result is very sensitive to the choice of the number of bins for the test. At the one end the χ^2 test allows one to use a single bin, which would not be rejected because the normalised number of samples will have the same area as the PDF. At the other end the bins can be chosen as fine as a single sample, that creates very large errors in the tails of the distributions. If the χ^2 test is to be used for measuring the accuracy of a PDF all the parameters have to be stated explicitly.

For this document the χ^2 test was sometimes performed on 100, or even 200 bins, and this has been found to be too many bins. This many bins allows for fine resolution in the body of the distribution, but this causes large errors in the tails. The errors in the tails can be reduced if larger bins are used but this reduces the resolution in the body. The advice from [30] was to merge bins to have at least 5 samples in a bin, but this has been found not to work well for the extremely long tails exhibited by some of the distributions analysed in the experiments. Sometimes single samples were separated by values of as much as 10 to 20. An improvement on the advice to merge the bins, could be to make the number of samples sum up to be dependent on the number of samples generated. Although it has not been investigated it might also be beneficial to investigate an approach where the bins are divided according to the theoretical percentiles, chosen to the required accuracy (say 1, 2 or 5 percent steps), because the percentiles will automatically scale the bin sizes with the shape of the distribution. For one million samples as was used for the experiments in this project, a 1% percentile division would result in ten thousand samples in a bin. For many cases this would not produce the fidelity required to accurately evaluate the tails. Further investigation into an approach to solve this problem is still required.

The addition of confidence intervals to the χ^2 test did provide better insight into the behaviour of the generated distributions. Confidence intervals meant that although the mean value seemed stable with change in distribution parameters, an increased variation in the results could be observed from the trends of the confidence intervals. Larger confidence intervals relate to more variation in the results, while narrow confidence intervals means there is little variation. A value of 99% confidence was chosen, which was too strict. A value of 95% would have sufficed, since actual measured clutter is very rarely an exact match to any model, be that model Weibull, Log-normal or Rayleigh distributed.

The KL divergence value was implemented as a statistical distance measure between the analytical and generated distributions. The actual size of the KL divergence value in itself did not provide meaningful insight. There is no indication if a value of 0.1 is close match to a distribution, or if it is a poor match. The KL divergence value however becomes useful when comparing it to other previously calculated values. This allows one to compare how well a certain aspect of a distribution is performing relative to itself, for example when the scale parameter is varied, or the correlation is increased. The KL divergence value is not perfect, since it ignores places in the measured histogram where no samples are generated, which causes unrealiable results. This was noticed for the Log-Normal distribution, where the KL divergence value showed improved performance for larger tails, even though the χ^2 showed otherwise, and the sampling in the tails became very sparse.

Measuring the higher order moments of the distributions and comparing them to the theoretical ones provided additional information that was not present when considering the results from the χ^2 tests alone. Although the measurement of all six moments from 1 to 6 seemed unnecessary because they appear to follow scaled versions of the same trend, it would still be useful to include them all in future research since they can alert one to other problems which were not encountered during this project. The higher the order of the moment the more relevant it is to the tail of the distribution, and this moment often signalled a lack samples in the tail region of the generated distributions. The degradation of the higher order moments were also more gradual than was observed with the χ^2 test, which is more in line with what one would expect.

Considering all the information presented in this document, many relevant results were obtained. The importance of well chosen PDF shape tests has been discussed and the relevance of the higher order moments was confirmed. This project highlighted certain problems with clutter simulation methods and proposed an alternative method to solve those problems. Although the alternate method introduced other problems, the method is not thoroughly understood yet and will require additional research to improve its performance, it exhibited relatively low complexity and good performance. The method is not thoroughly understood yet and will require additional research to improve its performance. The alternate method is more suitable for FPGA implementation, and since hardware in the loop radar environment simulation is the eventual goal for the work, this method might be more useful than those techniques in literature which focus purely on mathematical tractability and software simulation.

6 REFERENCES

[1] J. Davis, "Simulation Based Acquisition," Defence Modeling, Simulation and Tactical Technology Information Analysis Center, Huntsville, Alabama, Technical Report DMSTTIAC SOAR 99-02, April 1999.

[2] D. K. Barton, *Modern radar system analysis,* Artech House, 1988.

[3] N. Levanon, *Radar Principles*, First Edition, Wiley-Interscience, 1988.

[4] F. E. Nathanson, *Radar Design Principles*, Second Edition, SciTech Publishing, Inc., 1999.

[5] M. W. Long, *Radar Reflectivity of Land and Sea*, First Edition, Lexington Books, 1975.

[6] J. B. Billingsley, *Low-angle radar land clutter*. IET, 2002.

[7] F. T. Ulaby and M. C. Dobson, *Handbook of Radar Scattering Statistics for Terrain*, Artech House, 1989.

[8] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, Second Edition. McGraw-Hill, 1984.

[9] W. L. Simkins, V.C. Vannicola, and J.P. Ryan, "*SEEK IGLOO Radar Clutter Study*," Rome Air Development Center, Report No. Rept. TR-77-338 (DDC AD-A047 897), October 1977.

[10] D. K. Barton, *Radar System Analysis*, First Edition. Dedham: Artech House Inc., 1976, pp. 11-12.

[11] F. E. Nathanson, *Radar Design Principles*, First Edition. New-York: McGraw-Hill, 1969.

[12] K. D. Ward, R. J. A. Tough, and S. Watts, *Sea Clutter: Scattering, the K Distribution And Radar Performance,* Institution for Engineering and Technology, 2006.

[13] M. I. Skolnik, *Radar Handbook*, Second Edition. New York: McGraw-Hill, 1990.

[14] B. Binllingsley, A. Farina, F. Gini, M. V. Greco and P. Lombardo, "Impact of Experimentally Measured Doppler Spectrum of Ground Clutter on MTI and STAP," *Proceedings of 1997 International Radar Conference*, pp. 290-294, October 1997.

[15] P. Lombardo, M. Greco, F. Gini, A. Farina and J. B. Billingsley, "Impact of Clutter Spectra on Radar Performance Prediction," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 37, no. 3, July 2001.

[16] S. Sayama and H. Sekine, "Weibull, log-Weibull and K-distributed Ground Clutter Modeling Analyzed by AIC," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 37, 2001, pp. 1108-1113.

[17] M. Sekine and Y. Mao, *Weibull Radar Clutter*, IET, 1990, pp. 4-44, 67.

[18] P. Tsakalides and C. L. Nikias, "Robust Space-Time Adaptive Processing (STAP) in Non-Gaussian Clutter Environments," *IEE Proceedings - Radar, Sonar and Navigation*, vol. 146, issue 2, April 1999, pp. 84-93.

[19] E. Conte, M. Longo and M. Lops, "Modelling and simulation of non-Rayleigh radar clutter," *IEE Proceedings-F*, vol. 138, no. 2, April 1991.

[20] M. A. Richards, J. A. Scheer, and W. A. Holm, *Principles of Modern Radar: Basic Principles,* SciTech Publishing, Inc., 2010.

[21] M. I. Skolnik, *Radar handbook*, Third Edition, McGraw-Hill, 2007.

[22] D. B. Thomas, W. Luk, P. H. W. Leong, and J. D. Villasenor, "Gaussian random number generators," *ACM Computing Surveys (CSUR)*, vol. 39, Nov. 2007.

[23] R. C. C. Cheung, D.-U. Lee, W. Luk, J. D. Villasenor, "Hardware Generation of Arbitrary Random Number Distributions From Uniform Distributions Via the Inversion Method," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 15, no. 8, pp. 952-962, August 2007.

[24] D. B. Thomas and W. Luk, "Non-uniform random number generation through piecewise linear approximations," *IET Computers & Digital Techniques*, vol. 1, pp. 312-321, 2007.

[25] G. Marsaglia and W. W. Tsang, "The Monty Python Method for Generating Random Variables," *ACM Transactions on Mathematical Software*, vol. 24, no. 3, pp. 341-350, September 1998.

[26] J. G. Proakis and M. Salehi, *Communication Systems Engineering*, Second edition, Prentice Hall, 2002.

[27] S. Kullback, and R. A. Leibler, "On information and sufficiency," *Annals of Mathematical Statistics*, vol. 22, pp. 79-86, 1953.

[28] S. Kullback, "Letters to the editor," The American Statistician, vol. 41, no. 4, pp. 338-341, November 1987.

[29] D. C. Montgomery, G. C. Runger, and N. F. Hubele, *Engineering Statistics*, Third edition, Wiley, 2003.

[30] H. Chernoff, "The Use of Maximum Likelihood Estimates in χ^2 Tests for Goodness of Fit," *The Annals of Mathematical Statistics*, vol. 25, no. 3, pp. 579-586, Sep. 1954.

[31] R. A. Johnson, *"Miller & Freund's Probability and Statistics for Engineers,"* Seventh Edition, Upper Saddle River: Pearson Prentice Hall, 2005, p. 121.

[32] B. G. Lindsay and P. Basak, "Moments Determine the Tail of a Distribution (But Not Much Else)," *The American Statistician*, Vol. 54, No. 4, November 2000.

[33] F. Gini, M. Greco, M. Diani, and L. Verrazzani, "Performance Analysis of Two Adaptive Radar Detectors against Non-Gaussian Real Sea Clutter Data," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 36, No. 4, October 2000, pp. 1429-1439.

[34] M. Sankaran, "On the Non-Central Chi-Square Distribution," *Biometrika*, vol. 46, no. 1/2, pp. 235-237, Jun. 1959.

[35] M. Sankaran, "Approximations to the Non-Central Chi-Square Distribution," *Biometrika*, vol. 50, no. 1/2, pp. 199-204, Jun. 1963.

[36] W.C. Jakes, *Microwave Mobile Communications*, IEEE Press Reissue, 1994.

[37] J.E. Cilliers, "Multi-dimentional Lattice Equaliser for Q^2 PSK," Master's Dissertation, University of Pretoria, 2002.

[38] A. Farina, A. Russo, F. A. Studer and F. Scannapieco, "Reply on a Correspondence on a Coherent Model for Log-Normal Clutter," *IEE Proceedings,* vol. 134, no. 2, Apr. 1987.

[39] A. Farina, A. Russo and F. A. Studer, "Coherent Radar Detection in Log-Normal Clutter," *IEE Proceedings*, vol. 133, no. 1, Feb 1986.

[40] E. Conte and M. Longo, "Correspondence on a Coherent Model for Log-Normal Clutter," *IEE Proceedings,* vol. 134, no. 2, Apr. 1987.

[41] M. Rangaswamy, D. Weiner, and A. Öztürk, "Computer Generation of Correlated Non-Gaussian Radar Clutter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 31, no. 1, Jan 1995.

[42] G. Li and K. Lu, "Modelling and Simulation of Coherent Weibull Clutter," *IEE Proceedings*, vol. 136, no. 1, pp. 2-12, Feb. 1989.

[43] R. C. C. Cheung, D.-U. Lee, W. Luk, J. D. Villasenor, "Hardware Generation of Arbitrary Random Number Distributions From Uniform Distributions Via the Inversion Method," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 15, no. 8, pp. 952-962, Aug. 2007.