CHALMERS
UNIVERSITY OF TECHNOLOGY

Time-and-frequency synchronization for one-bit MU-MIMO

Feasibility and performance study

Master's thesis in Communications Engineering

CARL LINDQUIST

Department of Electrical Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2019

Master's thesis

Time-and-frequency synchronization for one-bit MU-MIMO

Feasibility and performance study

CARL LINDQUIST

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Electrical Engineering
Division of Communication and Antenna Systems
Chalmers University of Technology
Gothenburg, Sweden 2019

Time-and-frequency synchronization for one-bit MU-MIMO
Feasibility and performance study
CARL LINDQUIST

Supervisor: Sven Jacobsson, Ericsson Research and Chalmers University of Technology
Examiner: Giuseppe Durisi, Chalmers University of Technology

Master's Thesis
Department of Electrical Engineering
Division of Communication and Antenna Systems
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Typeset in LaTeX
Gothenburg, Sweden 2019

Time and Frequency Synchronization for One-bit MU-MIMO
Feasibility and performance study
CARL LINDQUIST
Department of Electrical Engineering
Chalmers University of Technology

Abstract

In this thesis, an uplink multiuser multiple-input multiple-output (MU-MIMO) system is considered. In the system, a number of user equipments (UEs) are communicating with a single base station (BS). The BS is fitted with several antennas and in each receiver chain, one-bit analog-to-digital converters (ADC) are used to convert the incoming signal from the analog to the digital domain. At each BS antenna, there will be an uncertainty regarding the timing as well as the carrier frequency of the received signal, typically referred to as symbol timing offset and carrier frequency offset. This thesis investigates the effect of these offsets on the communication system and provides an analytic expression for the signal-to-interference-noise-and-distortion ratio (SINDR). Moreover, an overview into the topic of synchronization itself is provided and some standard synchronization methods are described and evaluated in the context of one-bit MU-MIMO. The thesis demonstrates that despite the nonlinear distortion introduced by the one-bit ADCs, the system can still be synchronized. Lastly, the overall system performance in the presence of synchronization errors is discussed, as well as some ideas for future research.

Keywords: quantization, one-bit analog-to-digital converter (ADC), multiuser multiple-input multiple-output (MU-MIMO), symbol timing offset (STO), carrier frequency offset (CFO), synchronization

Acknowledgements

This thesis concludes my studies at Chalmers University of Technology, where I have been enrolled in the Communication Engineering Master's programme as the final part for the degree of Master of Science in Electrical Engineering. I am most grateful to everyone that has provided me with the necessary tools to complete this degree and this thesis.

The thesis has been carried out in collaboration with Ericsson Research at Lindholmen and I would like thank everyone that has contributed to my thesis in any way. In particular, I would like to extend a major thank you to my main supervisor, Sven Jacobsson at Chalmers University of Technology and Ericsson Research, for countless hours of assistance with this topic. I would also like to thank Dr. Keerthi Nagalapur for your patience with my questions and as well as valuable input.

Carl Lindquist, Gothenburg, April 2019

Contents

List of Figures

1

Introduction

Before the advent of digital communication, the way in which information was communicated was strongly influenced by the information *type*. Written words were typically transmitted as letters and the spoken word could for example be transmitted via telephone wires. These two systems both transported information, either via mailboxes, sorting systems and mail carriers or rotary dials, switchboards and copper wires but clearly, the requirements of the two are vastly different. This alludes to perhaps the main disadvantage inherent analog information transfer, namely the lack of flexibility. The information type, such as a voltage level or letter, dictates the requirements of the transmission channel. A phone call can not be transmitted via the postal system and vice versa. Since the second generation of mobile networks in the 1990's [1], the underpinnings of communication systems have been digital technologies. These systems need not concern themselves with the type of data to transmit — information in binary form is transmitted in the same fashion regardless of what the bits represent. This allows us to focus on developing strategies and techniques for data transmission in general, instead of having parallel development tracks for different applications, With that said, however, certain performance requirements, such as bit error rate (BER), data rate, spectral efficiency and latency of the system, will vary with the type of information the systems transmits. These are some of the standard performance metrics with which we benchmark systems and in terms of mobile communications, these are commonly used to set the specification of new generations.

Mobile communication networks has transitioned to a new generation roughly every tenth year since the first 1G network appeared in Japan [2] in the late 1970s. Largely adhering to this rule, the commercial release date for the fifth generation (5G) mobile network is imminent [3]. In 5G, a thousandfold increase in data rates, reduced end-to-end latency and supporting a ten- to a hundredfold increase in connected devices, are some of its targets [4, 5]. The first 5G networks have not yet been rolled out commercially, but numerous demonstrations have taken place. For example, KT Corp. showcased a number of applications of 5G during the 2018 Winter Olympic Games in South Korea, including live-streaming 360° video of competing athletes [6]. In May 2017, Ericsson together with Verizon also demonstrated live-streaming of 360° video. In the trial, a car was driving around a race track while the driver's only visuals where streamed to a set of virtual reality goggles from a camera mounted on the hood of the car [7].

To achieve the targets set for 5G networks, a number of technologies, both improvements of existing technologies as well as more radical ones, will need be deployed [8]. Two of the new technologies that will be incorporated in 5G are massive multiple-input multiple-output (MIMO) and millimeter-wave (mmWave). These technologies are considered key [5] and the work in this thesis is primarily related to these two. The next section provides a brief introduction to these topics.

1.1 Spectral congestion and mmWave

In the pursuit of higher data rates, we are intrinsically bound by the inverse relationship of symbol time and the bandwidth. This can be intuitively understood from the fact that if the symbol time is shortened, the signal will change more frequently, yielding a larger bandwidth. Whilst for example higher order modulation formats can be used to increase the rates, increasing the system bandwidth have been a trend throughout the development of mobile communication standards.

Increasing the bandwidth can lead to several challenges. First and foremost, the radio frequency spectrum, in which all mobile communication takes place, is a limited resource. As such, increasing the bandwidth indiscriminately is usually not an option. Below 6 GHz, where most wireless communication takes place in current mobile communication systems, the spectrum is extremely crowded, serving as a motivation to develop communication technologies outside of this band. One example of this is Visible Light Communication (VLC) [9], which, as implied by the name, uses visible light (that is, in the frequency band from 400–800 THz) as a carrier and overlays the light with a data-carrying signal. The actual communication is not carried out in the THz-domain; the visible light is merely used as a carrier wave. This technology is as of yet not particularly fast, but serves as an example of a means of wireless communication outside of the conventional band.

While auxiliary technologies such as VLC might slightly unload the mobile communication network for specific tasks, it will most likely never be the main workhorse of human communication. The majority of research attention is given to the mobile communication network itself and going forward, the move to higher frequencies is set to play a key part of the next mobile communications standard. Communications will take place both in bands below 6 GHz and in a number of different bands above 6 GHz [10], with the high frequency range in the 5G standard set to 24.25–52.6 GHz and the highest band allocated as of yet set to 37–40 GHz [11]. These higher frequency band are commonly referred to as mmWave, as the wavelength for frequencies ranging from 30–300 GHz have wavelengths ranging from 10–1 mm. Shorter wavelength has a number of effects, perhaps the most prominent being the increased susceptibility to blockages, due to shorter penetration depth. Further, as stated in the well-known Friis transmission equation [12], the received power of a signal transmitted in free space is inversely proportional to the square of its frequency, meaning that even without blockages, higher frequencies lead to less received power for a given antenna aperture.

One way to mitigate these issues is to use a technique known as beamforming. By using more than one antenna, we can adjust the shape of the beam according to some design criterion. Commonly, the beam is designed to interfere either constructively or destructively at one or more locations. Next, we will take a look at Massive MIMO, the technology underpinning beamforming strategies.

1.2 Massive MIMO

It is well-known that using more than one antenna can yield considerable performance improvements. Intuitively, this can be understood by considering a simple system where a message is transmitted to a single antenna through a noisy environment. The noise might have a very strong influence on some parts of the message, making it impossible for the receiver to completely recover what was transmitted. If, however, the receiver uses more than one antenna to pick up the message, the receiver would then have multiple versions of the same message. Now, if parts of the message strongly affected by noise in one version are unaffected in others, the receiver could then combine all versions of the message to hopefully recover the full message. A system where the receiver uses more than one antenna to pick up a signal from a single source, is known as single-input multiple-output (SIMO). MIMO refers to the case where both the transmitter as well as the receiver are equipped with more than one antenna. If designed properly, a MIMO system enable increased throughput, as it can establish parallel communication streams by adding the spatial dimension as a scheduling resource. A scenario where a single user is equipped with several antennas is sometimes referred to as single user (SU) MIMO and further, we can also consider many single-antenna users, known as multi user (MU) MIMO. Both SU-MIMO and MU-MIMO have been a part of both the Long Term Evolution (LTE) and Worldwide Interoperability for Microwave Access (WiMAX) standard for roughly ten years and has played a significant part in reaching target data rates [13].

Massive MIMO, as implied by the name, means that the number of antennas is high. Typically, the number of antennas at the base station (BS) is significantly higher than the number of antennas at the user equipments (UE). A key technique in Massive MIMO is beamforming, where the many antennas facilitates accurate direction of beams to discrete points in space. Massive MIMO has been an active research topic for several years and is expected to become an integral part of the next generation of mobile systems. Some theoretical benefits of Massive MIMO are capacity gains, increased robustness, and highly improved energy efficiency [14, 15]. Whether or not all of these theoretical promises can be realized at a reasonable cost remains to be seen, but regardless, Massive MIMO is a highly promising technology. The major interest in bringing this to market has spawned a number of new research areas related to the challenges associated with having large antenna arrays. In the following section, we will outline some of these challenges and provide the necessary motivation for this work.

1.3 Challenges

Having a large antenna array intrinsically entails an increased number of hardware components. Each antenna will require a hardware chain and for the elements to cooperate in some fashion, additional controlling units are needed. To keep the power consumption from skyrocketing, considerable research interest has been devoted to simplifying some hardware components. Examples of such, include the analog-to-digital converters (ADC) and digital-to-analog converters (DAC). As implied by their name, these devices are responsible for converting the output of a digital signal processor (DSP) to signal that can be fed to an antenna, or vice-versa. In the case of a ADC, these work by taking a snapshot of the analog signal at discrete timing instants and storing the amplitude of signal as closely as allowed by the resolution of the ADC. More precisely, an ADC performs both sampling (time-discretizing) and quantization (amplitude-discretizing). The discretizing operations are necessary as a digital device has finite memory, meaning the analog values can not be stored with infinite-precision. This thesis will focus primarily on the quantizing part of the ADC and the sampled and nonquantized signal will be commonly be referred to as the infinite-precision case.

Mathematically, we express the sampling part of the ADC as $y_s[n] = y_c(nT_s)$, where y_s is the sampled version of the continuous signal y_c at times $nT_s,\ n \in \mathbb{Z}^{\geq 0}$. The sampling time T_s specifies the time between two successive samples. The inverse of the sampling time $F_s = 1/T_s$ is called the sampling frequency. Next, the quantization operation can be described as $r[n] = \mathcal{Q}(y_s[n])$, where $r[n]$ is the quantized version of the discrete signal $y_s[n]$, according the- rules specified by the quantizer $\mathcal{Q}$. The number of possible outputs from the quantizer is known as the number of quantization levels. The number of quantization levels L is linked to the resolution of an ADC as $L = 2^b$, where b is the resolution in bits. For example, a three-bit quantizer has $L = 2^3 = 8$ possible output values, represented digitally as every permutation of a three-bit binary number.

In Figure 1.1, a number of examples of the sampling speed F_s and resolution b are shown. Beginning in the top left corner, the time-continuous signal y_c has been sampled with some speed F_s and then quantized with a three-bit quantizer. The red dots depicts the digital representation $r[n]$ of the signal and the quantization error is defined as $y_s[n] - r[n]$, i.e. the difference between the nonquantized and the quantized signal in the sampling instants. As the number of quantization levels is limited to $2^3 = 8$ levels, the quantization error is clearly discernible. In Figure 1.1, we have also added a zero-order hold reconstruction of the digital values, represented by the solid black line. This is a model of how the analog signal can be reconstructed from the digital values using a DAC that simply holds the digital value until the next sampling instant, producing a square-like output. Clearly, a using a first-order hold filter that linearly interpolates between successive sampling points would have produced an output that looks more similar to the original, but the purpose here is to demonstrate the effect of increasing the sampling speed. This is done in Figure 1.1b, where the sampling speed is $3F_s$. Comparing the quantization error in Figure 1.1a and 1.1b, we see that the magnitude of error is comparable, but the reconstructed line is much closer to the original. Though not evident from Figure 1.1b, there are additional benefits to increasing the sampling rate besides in the reconstruction phase, which will be mentioned later.

In Figure 1.1c, the sampling speed is again set to F_s and the resolution increased to 10 bits. With $b = 10$, the number of quantization levels increases to $2^{10} = 1024$, yielding a quantization error that is virtually zero.

In Figure 1.1a, 1.1b and 1.1c, we have used what is known as an automatic gain control (AGC). This is a device which attenuates the input signal so that it fits within the range of the quantizer levels. Defining q to be the difference between two adjacent quantization levels, this implementation of the AGC ensures that the maximum input value is $q/2$ above the highest quantization level. Then, the maximum size of the error will be $q/2$, regardless of the signal level. Without the AGC, the quantizer would have had to select its highest or lowest value as soon as the input went outside its

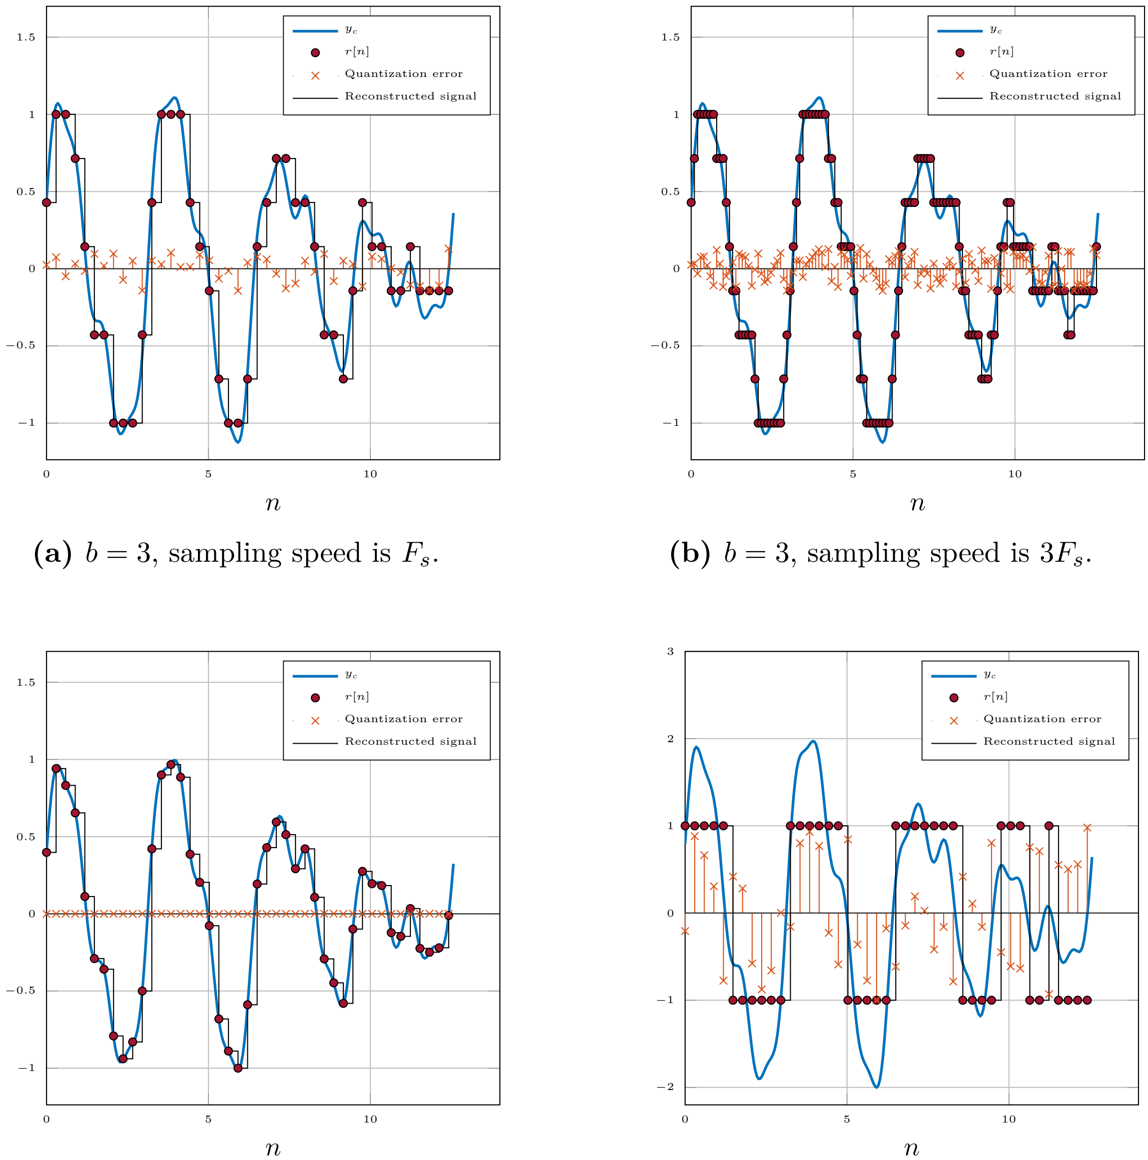

(a) $b = 3$, sampling speed is F_s.

(b) $b = 3$, sampling speed is $3F_s$.

(c) $b = 10$, sampling speed is F_s.

(d) $b = 1$, sampling speed is F_s.

Figure 1.1: Some quantizers with different parameters.

range, yielding significant errors at the extremes. This is an example of saturation (or clipping) and is a well-known effect in for example amplifiers, arising when the dynamic range of the input exceeds the dynamic range of the output.

Lastly, in Figure 1.1d, the resolution is set to a single bit. Note that having an AGC or not does not make a difference, as only the sign of the incoming signal is stored; the magnitude is of no importance. Consequently, we can remove the AGC from the receiver chain, which is one of the reasons why one-bit quantizers have generated significant research interest in the last few years. Fascinatingly, and perhaps rather counter-intuitively, we will see that the very apparent quantization error in Figure 1.1d is not enough to render a MIMO system unusable. In fact, research has demonstrated quite the opposite, as studies has demonstrated not only that systems using low-resolution converters does in fact work — only a few bits are required to get similar performance to that of the infinite-precision case [16].

As is clear from Figure 1.1, in order to improve the performance of a ADC, we can either sample the signal more often or store the value with greater precision. Unfortunately, increasing the sampling rate and resolution of an ADC also increases the power consumption. For a given sampling rate F_s Hz and resolution b bits, the power dissipation scales as roughly [17]

$$P \propto 2^{b_{\text{eff}}} F_s \tag{1.1}$$

where b_{eff} is the effective number of bits. It is a standard way of measuring the true resolution of an ADC, taking into account the distortion introduced by the ADC circuit itself.

From (1.1), we see that for each additional bit of resolution, the power consumption scales by roughly a factor two. Moreover, in order to compare different architectures and identify trends, a figure of merit (FOM) was proposed by Walden in [18] as

$$F_{\text{W}} = \frac{2^{b_{\text{eff}}} F_s}{P} \tag{1.2}$$

This FOM, now referred to as Walden's FOM, is in commonly in use and in Figure 1.2, the inverse of (1.2) with each data point divided with its resolution is shown. The envelope have been constructed from the average of the five best (with relation to their respective FOM) designs [19].

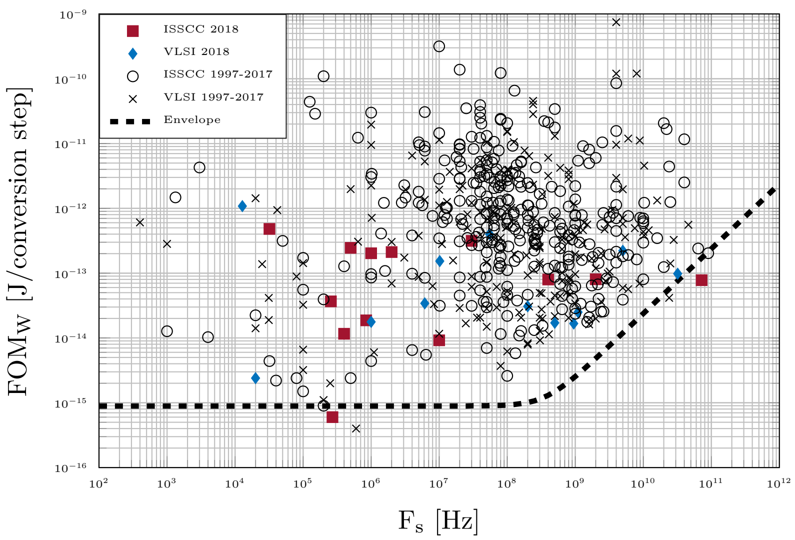

Figure 1.2: Walden's Figure of Merit. Data gathered from the International Solid-State Circuits Conference (ISSCC) and the Very Large Scale Integration Symposium (VLSI). Data fetched from [19].

We see that increasing the speed of an ADC can indeed be costly in terms of power consumption. High speed converters can bring a number of benefits, such as radio frequency (RF) sampling, where the received signal is sampled and converted into the digital domain at passband directly, removing a number of analog components required in the mixing stage [20]. This reduces complexity and cost, as well as decreases power consumption. Further, it has been shown for example in [21] that oversampling can yield performance benefits if low-resolution converters are used.

With future mobile communication network likely to operate at wider bandwidths than today, it seems likely that the demand for fast ADCs is not going to diminish. Consequently, reducing the resolution of the converters could be a solution. For example, a typical ADC deployed in a BS today has a resolution of 10–14 bits [22]. Replacing those converters with ADCs and DACs with lower resolution has the potential to yield significant energy savings, at least when examining the converters in isolation. As shown in Figure 1.1, the quantization error grows as the resolution is decreased, introducing more distortion to the system. Adding more antennas, as is the core concept in Massive MIMO, will provide array gain that to some degree mitigates the effect of increased quantization distortion, while at the same time increasing the total power consumption. This thesis exclusively compare the infinite-resolution case and the extreme one-bit quantized case and for a more in-depth discussion into this topic of low-resolution converters and their role in future Massive MIMO systems, see for example [23].

For one-bit ADCs to be a viable technology in future communication systems, it needs to be shown that acceptable performance can still be achieved despite the major impairments that they bring. Research so far has demonstrated the viability of low-resolution quantizers in a number of aspects, indicating that this indeed a promising path forward. A topic currently receiving some attention with promising results [24, 25, 26, 27] concerns the synchronization of systems employing low-resolution converters. Demonstrating that systems employing one-bit ADCs can still be synchronized is an important step in the research of this technology and the main motivation for this thesis.

1.4 Outline of thesis

In this thesis, some aspects of one-bit converters will be investigated. Specifically, it will be examined whether systems employing these types of converters can be synchronized using standard methods. In Chapter 2, the system model and channel input-output will be discussed. along with an overview of the effects of imperfect synchronization. In Chapter 3, the necessary mathematical tools will be mentioned and a more in-depth analysis of the effects mentioned in Chapter 2 will be performed. Expressions for the power in the received signal and signal-to-interference-noise-and-distortion-ratio (SINDR) will also be derived. In Chapter 4, some well-known synchronization strategies will be discussed and their performance in the non-quantized and the one-bit quantized case will be compared. In Chapter 5, the overall performance of the system with synchronization errors present will be examined via simulations and lastly, in Chapter 6, the results from this thesis will be discussed. Some ideas for future research into this field will also be suggested.

2

System model

This chapter explains the system model, the assumptions made regarding the communication system and the effects of imperfect synchronization. Here, the aim is to give a complete characterization of the problem at hand, before delving into the analysis.

2.1 Orthogonal Frequency Division Multiplexing

Before going into the channel input-output model, it is fitting to say something with regards to the modulation format used in modern communication systems. The medium through which the signal is transmitted is called the channel and can be characterized by a number of parameters. Noise introduced by for example hardware, is usually modeled as an additive effect. As this in intended to capture a wide array of noise sources into a single additive effect, it can by the central limit theorem be approximated as a Gaussian random variable. In addition to this, wireless channels also have a multiplicative effect known as fading. A natural distinction between different fading environments is whether the channel has the same effect on all frequencies in the signal bandwidth or if its frequency response varies. The former is usually referred to as a frequency-flat channel, arising in narrow-band scenarios where the communication bandwidth is smaller than the channel coherence bandwidth. The latter case is referred to as a frequency-selective channel and usually arises in environments where the reflection and refraction of a transmitted signal must be taken into account.

With growing signal rates, the system will occupy a larger bandwidth, introducing additional design concerns. In a single-user scenario, no other user is competing for the same frequency resource, but if, however, multiple users are present, some sort of scheduling in order to prevent collisions between messages must be implemented. One solution would be to designate a slot in time for each user, in which the user is granted access to the channel. This idea is the basis of time-division multiple access (TDMA) and was the main multiple access technique in early mobile systems. It is still used today, for example in LTE where it is used in conjunction with other techniques [28].

While TDMA indeed solves the issue of multiple users simultaneously accessing the channel, but does nothing to alleviate another issue, namely the frequency-selectiveness of the channel. If the signal bandwidth is larger than the channel coherence bandwidth, it means that the spectral gain is not constant over the entire signal bandwidth, making equalizing more complex. A response to this is to use orthogonal frequency division multiplexing (OFDM), which converts a single wideband channel into several narrowband channels with lower rates. Dividing the available frequency spectrum into several channels, known as a multi-carrier modulation, had been around for a while, but the idea of modulating each symbol with signals from an orthogonal set is normally credited to [29]. Using an orthogonal set of signals, we are able overlap the frequency bands without incurring any inter-carrier interference (ICI), thus maximizing the transmission rate. Further, frequency resources can be dynamically allocated based on the current conditions of the sub-channels, thereby avoiding energy being wasted on unusable parts of the spectrum. The first major application of the format was the in asymmetric digital subscriber line (ADSL) in the 1990's [30]. By then, efficient implementation had been made available via fast DSP units and fast Fourier transforms (FFT). For a more complete description of the development of OFDM, see e.g. [31].

In the modulation process, the inverse discrete Fourier transform (IDFT) is used to modulate each data symbol $\hat{x}[k] \in \mathbb{C}$ with a complex exponential function drawn from an orthogonal set. Specifically, for N sub bands

$$x[n] = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} \hat{x}[k]\ e^{j2\pi nk/N}, \quad n = 0, 1, \ldots, N-1 \tag{2.1}$$

where $\hat{x}[k]$ is the data symbols drawn from some constellation and $x[n]$ the time-domain signal to be transmitted. The scaling factor $1/\sqrt{N}$ in front of the summation is there to ensure that the average power of the signal is not affected by the IDFT operation. The exponentials $\mathrm{e}(j2\pi nk/N), \quad \forall n \in \{0, 1, \ldots, N-1\}$ in (2.1) are commonly referred to as subcarriers and are what the data to be transmitted is modulated with. The phase difference of the complex exponential functions being multiplied with $\hat{x}$ is $1/N$ and to verify that the set spanned by these complex is orthogonal, note that

$$\begin{aligned}
&\sum_{n=0}^{N-1} e^{j2\pi nk/N} \cdot e^{-j2\pi np/N} \\
&= \sum_{n=0}^{N-1} e^{j2\pi \frac{k-p}{N} n} \\
&= 0\ ,\ \forall\ k \neq p.
\end{aligned}$$

Not all N values of $\hat{x}[k], k = 0, 1, \ldots, N-1$ needs to contain a data symbol. Some zeros are commonly inserted at specific points in order to increase the robustness of the system. This will however not be considered in this thesis and all subcarriers will be assumed to carry data symbols, which will be referred to as symbol-sampling rate.

In the demodulation process, the frequency-domain symbols are retrieved by computing the discrete Fourier transform (DFT) of the time-domain signal, defined as

$$\hat{x}[k] = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} x[n] \; e^{-j2\pi nk/N}, \quad k = 0, 1, \ldots, N-1 \tag{2.2}$$

where the scaling $1/\sqrt{N}$ again ensures that the average power after the DFT is unchanged.

In a fading channel, some concessions regarding the transmission must be made in order for the orthogonality between subcarriers to be preserved. Specifically, each OFDM symbol is commonly cyclically extended via the insertion of a cyclic prefix (CP). The CP consists of the last $P \leq N$ samples of the original OFDM symbol, bringing the total length of the OFDM symbol to $N + P$. Naturally, as some parts of the symbol are repeated, this redundancy lowers the total rate. Another issue with OFDM is the peak-to-average ratio (PAPR), which captures how much the amplitude of the time-domain signal varies. In OFDM, compared to a single-carrier system, this ratio tends to be fairly high, spelling difficulty for the amplifiers. Typically, an amplifier is only linear in a limited region, meaning that we are more likely to incur nonlinear distortion for signals with large amplitude variations. This thesis will assume that this effect, as well as a number of additional hardware-related concerns, can be disregarded. For more details on these matters, see e.g. [32].

2.2 Channel input-output model

In this work, we consider an uplink baseband model where U single-antenna UEs are communicating with a BS equipped with B antennas. Operations such as filtering, up and down conversion and mixing will be assumed to be ideal and therefore not explicitly considered.

An overview of the considered MU-MIMO OFDM system is depicted in Figure 2.1. The raw bits coming into our system are assumed to be uniformly distributed and fed to a modulator unit, generating U streams of Gray-coded complex symbols drawn from a quadrature phase-shift keying (QPSK) alphabet, i.e. $\hat{x}_u[k] \in \{\pm 1/\sqrt{2} \pm 1/\sqrt{2}j\}$ for $u \in \{1, 2, \ldots, U\}$. Next, the symbols $\hat{\mathbf{x}}[k] = [\hat{x}_1[k], \hat{x}_2[k], \ldots, \hat{x}_U[k]]^T \in \mathbb{C}^U$ are transformed into the time-domain via the IDFT defined in (2.1), a cyclic prefix is added and the time-domain signal $\mathbf{x}[n] = [x_1[n], x_2[n], \ldots, x_U[n]]^T \in \mathbb{C}^U$ is transmitted over a channel $\mathbf{H} \in \mathbb{C}^{B \times U}$. In this thesis, the channel will be assumed to be frequency-flat, i.e. the number of channel taps, modelling the delay spread, is assumed to be one. Upon arrival at the receiver, there will be some uncertainty regarding the absolute timing as well as the carrier frequency of the signal. These effects are called the symbol timing offset (STO) and carrier frequency offset (CFO), respectively, and are a major focus of this thesis.

At the receiver, the received signal $\mathbf{y}[n] = [y_1[n], y_2[n], \ldots, y_B[n]]^T \in \mathbb{C}^B$ is fed to $2B$ ADCs, one pair for each receiver chain, where the real and imaginary part of the signal at each antenna element will be quantized separately. The resolution is limited to one bit, effectively only storing the sign of the incoming signal. Next, the STO and CFO will be estimated and compensated via some method, examined in more detail in Chapter 4. If the estimation of the STO and CFO is not perfect, then there will be residual STO and CFO after the compensation stage, potentially affecting the performance of the system. The precise effect of the residual STO and CFO will be examined in Chapter 3.

After the STO and CFO estimation and compensation stage, the CP is removed and a DFT is performed. Lastly, to obtain an estimate $\hat{\mathbf{x}}^{\text{est}}[k] \in \mathbb{C}^U$ of the transmitted symbols $\hat{\mathbf{x}}[k]$, the signal is passed through a equalization stage. After equalization, the estimated symbols are then fed to the decision and decoding units of the system, finally yielding an estimate of the transmitted bits.

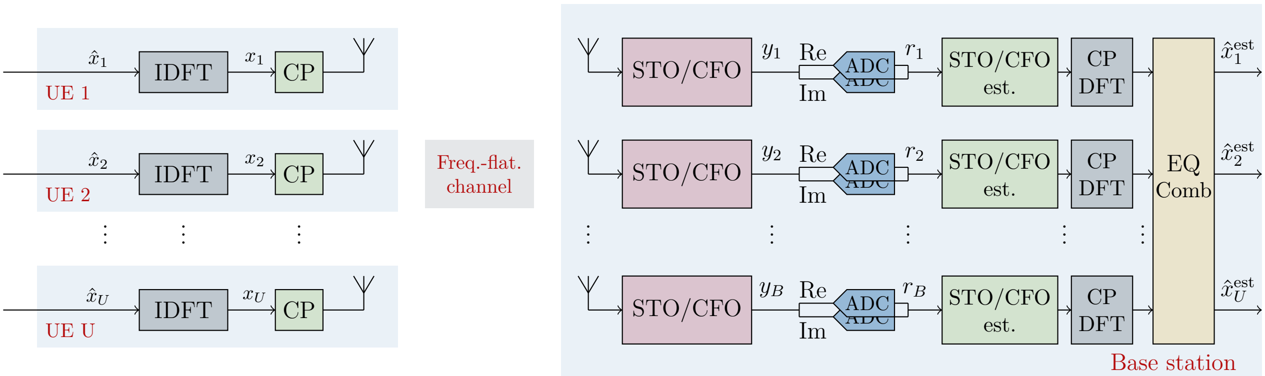

Figure 2.1: The system model for U single-antenna transmitting UEs and B receive antennas.

Assuming that the users are transmitting a continuous stream of OFDM symbols, the received signal at the bth antenna $y_b[n]$ can be expressed as

$$y_b[n] = \sum_{u=1}^{U} \mathrm{e}^{j2\pi\epsilon_{0,u}n/N}\; h_{b,u}\; x_u[n+\delta_{0,u}] + w_b[n]. \tag{2.3}$$

In (2.3), $\epsilon_{0,u}$ and $\delta_{0,u}$ represents the CFO and STO between the BS and the uth UE, respectively. The scalar value $h_{b,u}$ represents the channel from user u to antenna b and $x_u[n]$ is the transmitted signal from user u. The term $w_b[n]$ stands for the thermal noise at antenna b, $w_b[n] \sim \mathcal{CN}(0, N_0)$, where N_0 is the noise power spectral density.

The quantized signal $r_b[n]$ is defined as $r_b[n] = \mathcal{Q}\left(y_b[n]\right)$, where $\mathcal{Q}(\cdot)$ describes the nonlinear quantizer operation. It is defined as

$$\mathcal{Q}(y_b[n]) = \frac{1}{\sqrt{2}}\Big(\mathrm{sgn}\left(\mathrm{Re}\{y_b[n]\}\right) + j\ \mathrm{sgn}\left(\mathrm{Im}\{y_b[n]\}\right)\Big) \tag{2.4}$$

where $\mathrm{sgn}(\cdot)$ is the signum function. The scaling $\sqrt{1/2}$ is chosen so that $\mathbb{E}[|r_b[n]|^2] = 1$.

Lastly, an estimate of the transmitted symbols from user u, $\hat{x}_u^{\text{est}}[k]$, is obtained as

$$\hat{x}_u^{\text{est}}[k] = \sum_{b=1}^{B} a_{u,b} \hat{r}_b[k] \tag{2.5}$$

where $a_{u,b}$ is the (u, b)th entry in the equalization matrix $\mathbf{A} \in \mathbb{C}^{U \times B}$ and $\hat{r}_b[k]$ is the DFT of $r_b[n]$.

In this thesis, a zero-forcing (ZF) equalizer is considered, which attempts to invert the effect of the channel via the psuedo-inverse of the channel, i.e.

$$\mathbf{A} = \left(\mathbf{H}^H \mathbf{H}\right)^{-1} \mathbf{H}^H. \tag{2.6}$$

The analysis of the received signal is performed in Chapter 3.

We will use this model as the basis for our investigation and next, we will give an overview of the effects of STO, CFO, and one-bit quantization on the signal.

2.3 Impact of imperfect synchronization

This section will provide an overview of the adverse effects STO and CFO will have on the received signal. In a real-world system, there are additional synchronization-related matters to consider, such as the sampling clock offset, but we will limit ourselves to the STO and CFO. For a discussion on the effects of other synchronization offsets, see e.g. [33].

This section intends to provide the intuition for the effects of STO and CFO. For a more complete mathematical analysis, please refer to Section 3.

2.3.1 Symbol timing offset

Generally speaking, the receiver in a communication system does not know in advance when the transmitter will send something. This means that the designer of the system will need to devise a way for the receiver to automatically detect when a signal is present in the channel. This process is known as *frame detection* and in Section 4, some strategies for accomplishing frame detection are mentioned. Determining exactly when a transmitted message has arrived at the receiver is not trivial and uncertainty in that process can lead to STO, the topic of this section.

We will assume that frame detection and STO estimation has already been performed and that the timing uncertainty is in the order of $\pm N$, the length of the OFDM symbol. As mentioned in Section 2.2, this remaining STO will be referred to as residual STO and we define the residual STO from the uth user to be $\delta_u = \delta_{0,u} - \delta_u^{\text{est}}$. From Figure 2.1, the next step after synchronization offset estimation is the removal of the CP and a DFT operation. The continuous stream of received samples is divided into blocks of size N, which are fed to an N-point DFT. We refer to the ith block as the ith DFT window. This is illustrated in Figure 2.2. If $\delta_u = 0$, then the ith DFT window contains precisely the N samples from the ith symbol. Clearly, if there is no residual STO, no additional interference will be introduced.

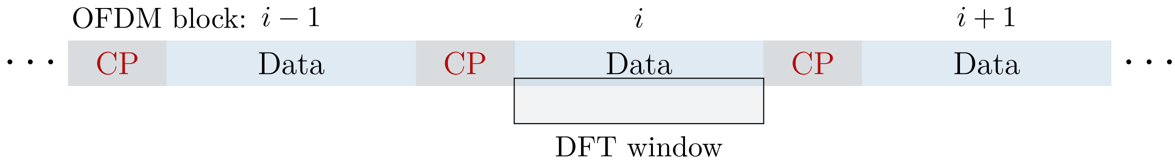

Figure 2.2: A depiction of the correct placement of the DFT window.

If, however, $\delta_u \neq 0$, the residual STO might cause interference in the system. There are several cases to consider, as the DFT window will either have its starting point inside the CP of the ith symbol or contain some samples from either the previous symbol $i-1$ or from following symbol $i+1$. These cases are denoted Case (i), Case (ii) and Case (iii), respectively, and in the last two cases, as samples belonging to a different OFDM symbol are included in the DFT window, inter-symbol interference (ISI) is introduced. Moreover, the misaligned DFT window also causes self-interference in the form of ICI.

In the following discussions, a single-user scenario transmitting over an additive white Gaussian noise (AWGN) channel with no quantization (i.e. infinite-resolution) is considered where the signal-to-noise ratio (SNR) is set very high, such that (2.3) simplifies to

$$y[n] = \mathrm{e}^{j2\pi\epsilon_0 n/N}\ x[n+\delta_0]. \tag{2.7}$$

The corresponding frequency-domain symbols is then found via the DFT, i.e. $\hat{y}[k] = \mathrm{DFT}\{y[n]\}$. The purpose of disregarding thermal noise and MU interference is to demonstrate only the effect of the STO. Since there is only a single user, the subscript u in δ_u and ϵ_u will be temporarily dropped and lastly, throughout this thesis, $\delta_u > 0$ will be taken to mean that the DFT window contains samples from the *next* symbol.

Case (i): $-P < \delta < 0$

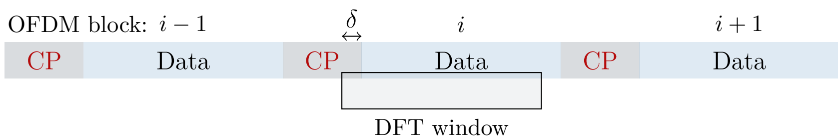

In this case, the starting point of DFT window is taken too early, but still within the part of the cyclic prefix. We will miss a few samples of the symbol in the end, but as the cyclic prefix contains the information that we missed, we will able to perfectly reconstruct the transmitted data. A cyclic shift in the time domain will appear as a linear phase shift for all subcarriers in the frequency domain, which can be rectified with an equalizer. The effect is depicted in Figure 2.3a.

We clearly see that there is no added ISI or ICI. In Figure 2.3, the axes are the in-phase (I) and the quadrature (Q) part of the signal.

Case (ii): $-N \leq \delta \leq -P$

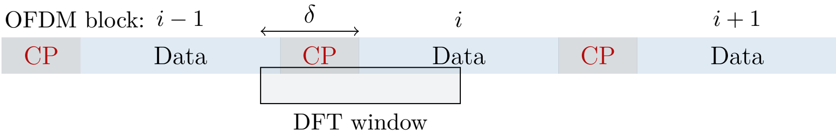

In this case, the DFT window will miss δ samples of symbol i and instead get δ samples from the $(i-1)$th symbol. The presence of erroneous samples in the DFT window will cause both ISI and ICI and since some samples of the desired part are missing, there will be an attenuation of the desired part of the signal. The misaligned DFT window will also cause a rotation of the received constellation, depicted in Figure 2.3b.

Case (iii): $0 < \delta \leq N$

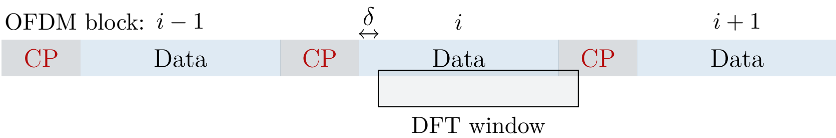

The effect in this case is in the frequency-flat channel model identical to Case (ii). The situation is depicted in Figure 2.3c, where $\delta > 0$. The received frequency-domain symbols are rotated and scaled and from the spreading of the symbols, it is evident that interference has been introduced.

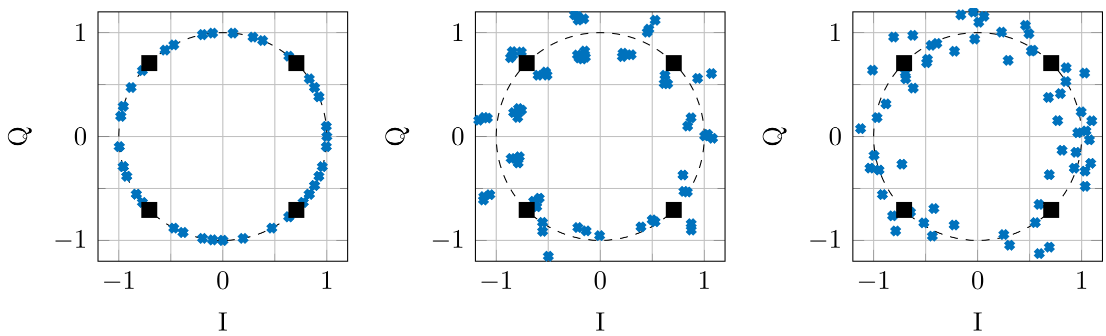

(a) Case (i), $\delta = -13$ **(b)** Case (ii), $\delta = -18$ **(c)** Case (iii), $\delta = 2$

Figure 2.3: The effect to a QPSK constellation in some cases of residual STO in a single-user AWGN channel with high SNR, $N = 64$, $P = 16$.

2.3.2 Carrier frequency offset

In order to properly decode a received message, the receiver will need to have accurate knowledge of the carrier frequency. The data is modulated onto a carrier wave with some frequency f_c and without knowledge of this parameter, the receiver will not be able to shift the signal down the baseband. Due in part to imperfect oscillators at the transmitter and receiver, there will be a difference between the reference clocks in either end. Additionally, if the receiver or transmitter is moving with respect to the other, the Doppler effect will cause a change to the frequency of the transmitted signal. The difference in the carrier frequency at the transmitter and receiver is called CFO. This can have a detrimental effect on our received signal if not mitigated, especially in systems employing OFDM, where exact knowledge of the center frequency in each sub band is key to preserve the orthogonality between subcarriers. In the context of OFDM, the CFO is commonly measured in relation to the subcarrier spacing, Δf, and we define the CFO associated with the uth UE as $\epsilon_{0,u} = (f_{c,u} - f'_{c,u})/\Delta f_u$, where $f'_{c,u}$ denotes the carrier frequency estimate in the receiver. As in the case of the STO, the CFO will be estimated and compensated before the DFT operation and we denote the residual CFO after estimation as $\epsilon_u = \epsilon_{0,u} - \epsilon_u^{\text{est}}$.

To explain the effect of different values for the CFO, we begin by defining the set of exponentials $\mathcal{A}_t = \{\exp(j2\pi nk/N) \ : \ k = 0, 1, \ldots, N-1\}$. These are each of the complex exponentials from (2.1) and represent the center frequencies of the sub bands in the transmitted signal.

Next, define another set of exponentials $\mathcal{A}_r = \{\exp(-j2\pi n(k+\epsilon_u)/N) \ : \ k = 0, 1, \ldots, N-1\}$. These are the center frequencies in transmitted signal, each affected by a phase rotation proportional to ϵ_u. These are what is used to demodulate the received signal and note that while the demodulated symbols will be rotated if $\epsilon_u \neq 0$, we see that if $\epsilon_u \in \mathbb{Z}$, then $\mathcal{A}_r = \mathcal{A}_r$. This means that as long as ϵ_u is integer-valued, the received signal will be demodulated with the same orthogonal set it was modulated with in the transmitter, however in a different order. Since orthogonality is preserved, no ICI is introduced and given that this integer-value for ϵ_u can be found, it will be possible to perfectly retrieve the transmitted symbols from the received signal. However, if $\epsilon_u \notin \mathbb{Z}$, then $\mathcal{A}_t \neq \mathcal{A}_r$. This introduces ICI and can potentially cause significant performance degradation.

Along the lines of the previous paragraph, it is common practice to divide the CFO into an *integer* part and an *fractional* part, as their respective impact on the received signal is very different. In the integer case, the phase difference between two consecutive entries in either $\mathcal{A}_t$ or $\mathcal{A}_r$ is precisely Δf and we can conclude that the effect of an integer CFO is that the output of the DFT is cyclically shifted with respect to the transmitted sequence. No ICI has been incurred and given that we can somehow acquire the integer part of the CFO, we will be able to perfectly reconstruct the transmitted sequence.

The fractional part of the CFO can, however, be potentially devastating. In the context of OFDM, it will cause every entry $a_r^k \in \mathcal{A}_r$ to differ from $a_t^k \in \mathcal{A}_t$, resulting in ICI. In Figure 2.4, we see the effect of CFO on our received OFDM signal. The SNR is again set very high, so the only visible effect is that of the CFO. We see that already at $\epsilon_u = 0.1$, in Figure 2.4b, a significant amount of noise has been added. At $\epsilon_u = 0.3$, in Figure 2.4c, the effect is so severe that even if ϵ_u was known and the induced rotation compensated for, the CFO will still make it unable to correctly decode every point. These figures clearly illustrate the importance of accurate CFO estimates and the effect will be more closely examined in Section 3.

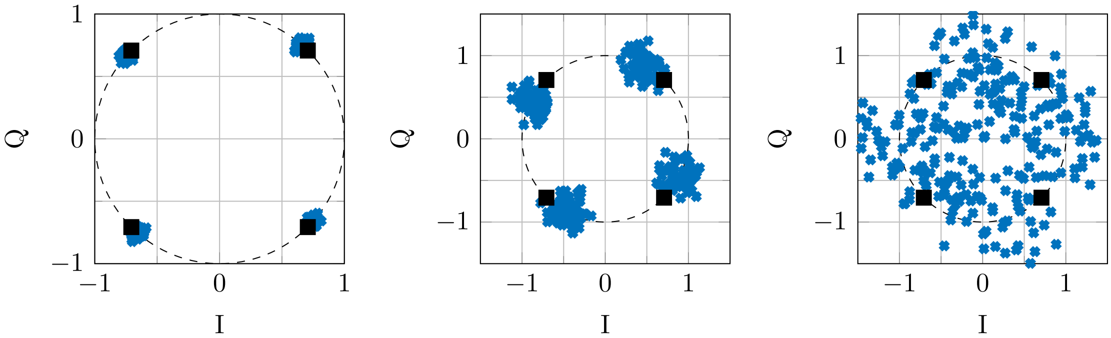

(a) $\epsilon = 0.02$ **(b)** $\epsilon = 0.1$ **(c)** $\epsilon = 0.3$

Figure 2.4: The effect of the CFO on an OFDM symbol with 256 subcarriers, single-user scenario with AWGN and high SNR.

2.3.3 One-bit quantizers

A one-bit quantizer is a device that will only store the sign of the incoming signal. We will quantize the real and imaginary part of the received baseband signal separately, as indicated in Figure 2.1. Obviously, a lot of information is lost in that process. Assuming a single user transmitting over an AWGN channel to a receiver consisting of only a one-bit quantizer, then, in every sampling instant, the output of the quantizers would, regardless of the input, be a member of the set $\mathcal{X} = \gamma\{\pm 1 \pm j\}$, where γ is some scaling factor. In a single-antenna system, we could not support any modulation format of higher order than QPSK, as there simply would not be any additional information to utilize. However, if each transmitted symbol were received on more than one antenna with independent noise realizations, then each received symbol could potentially be combined to gain additional insights about the transmitted symbol. The intuition for this can be obtained by examining Figure 2.5.

In Figure 2.5a and 2.5b, the clouds of blue dots represents a number of realizations of $\hat{a} + \hat{n}$, where $\hat{a}$ is the top-left constellation point in a 16-point quadrature amplitude modulation (16-QAM) constellation (marked with a black square) and $\hat{n} \sim \mathcal{CN}(0, 1/\text{SNR})$. The plots are zoomed in so that only the second quadrant is visible. The real and imaginary part of each point is then quantized separately with a one-bit quantizer described by (2.4). The quantization point is marked with the yellow diamond and lastly, all quantized points are averaged to form the orange dots.

In Figure 2.5a, we see that almost all of received symbols are in the second quadrant, and as hardly any received points have crossed into to any quadrant other than the second, almost all points are quantized to the same quantization point. Consequently, the average of the quantization points is very close the quantization point itself.

In Figure 2.5b, however, a significant number of the received points show up in the other three quadrants. The average of the quantized data is therefore shifted towards the origin, making it possible to distinguish between the transmitted points in Figure 2.5a and 2.5b.

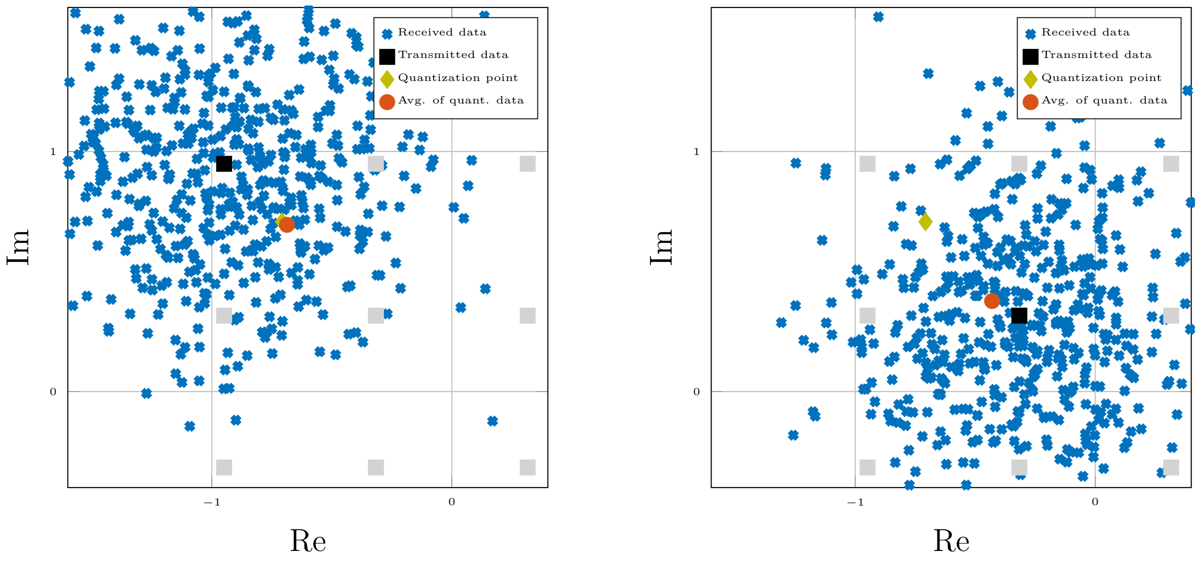

(a) **(b)**

Figure 2.5: Around 500 realizations of $\hat{a} + \hat{n}$, with SNR $= 5$ dB.

Many realizations of a received symbol with independent noise realizations approximately describes a multi-antenna system. Note that while the thermal noise on each antenna could plausibly be modelled as independent, interference sources are usually not uncorrelated over all antennas, so this description is not completely accurate. However, at least part of the noise associated with the signal on each receive antenna can be assumed to be independent, so even if correlated noise sources can not be completely averaged out, multi-antenna systems are able to support high order constellations, as shown in e.g. [34]. Note that the number of data points in Figure 2.5 is highly excessive when viewed as individual antenna elements - the illustration is only meant to clearly demonstrate why one-bit quantization supports higher-order modulation formats.

An observation from Figure 2.5, is that it would have been impossible to resolve the 16-QAM points if the SNR $\gg 1$, as the received symbols in Figure 2.5a and 2.5b would have been quantized to the same point. This is a perhaps surprising result that implies that for a given set-up, there will be an optimal level of thermal noise in terms of symbol error rate. This would also means that the system performance could benefit from the introduction of additional noise. Framed in the context of dithering, the intuition for this effect may become clear. Improving performance using dithering is a well-known technique in a number of fields, such as image or audio processing. It is not uncommon for dithering to be used as a mean to artificially increase the variability of a data set, which, as demonstrated in Figure 2.5, is critical for the average of the quantizer output to fall in set with higher cardinality than four.

Lastly, we note that the ability of one-bit quantizers to support higher order modulation formats in a multi-antenna setup, hints that OFDM can be used as a transmission scheme. In the time-domain, OFDM can be viewed as a higher-order modulation. As demonstrated, these can be supported by one-bit quantizers if the number of antennas is sufficiently high. This general conclusion is also valid for a frequency-selective channel, demonstrated in for example [16]. The topic of low-resolution quantizers and for reading on the topic, see for example [35, 21, 36, 37].

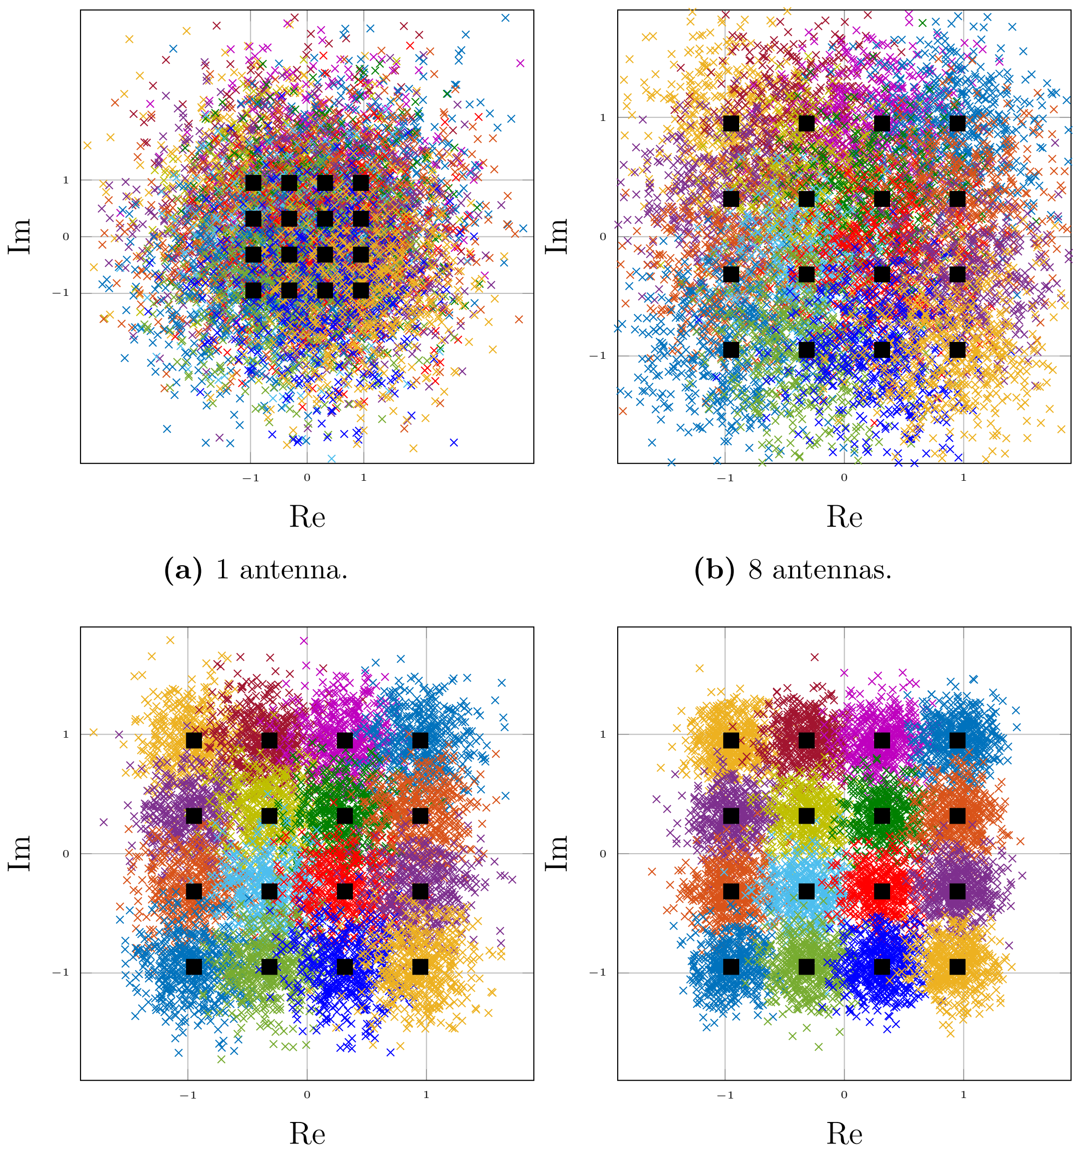

(a) 1 antenna. **(b)** 8 antennas. **(c)** 32 antennas. **(d)** 128 antennas.

Figure 2.6: 32 OFDM symbols with 256 subcarriers transmitted over AWGN and quantized and averaged at the receiver, SNR = 0 dB.

In Figure 2.6, we see the effect of adding antennas to the system. The constellation is gradually becoming more and more defined and it is interesting to note that even with a modest number of antennas, such as in Figure 2.6c, the constellation is clearly discernible. Given that the number of antennas in Massive MIMO is likely to be on the order in 64 or larger [15], we can conclude that one-bit quantizers could be a viable solution to reduce power consumption whilst retaining support for higher order constellations.

3

Analysis

Next, we will examine the effects of CFO and STO on a signal that is one-bit quantized. In order to successfully perform an analysis of the quantized signal, we will need a tool to handle this kind of nonlinear amplitude distortion. One such tool is *Bussgang's theorem*, which will be introduced and explained in Section 3.1. We will derive the closed-form expression for the received frequency-domain signal in the presence of STO and CFO and then use Bussgang's theorem to extend the analysis to the one-bit case.

3.1 Bussgang's theorem

A useful tools in the analysis of these kinds of system is Bussgang's theorem, named after Julian J. Bussgang who published it in 1952 [38]. In its original formulation, it states that for two Gaussian signals, their cross-correlation will be the same up to a scaling before and after one of the signals has undergone a nonlinear amplitude distortion. This situation is depicted in Figure 3.1a, where $x(t)$ and $y(t)$ are two (generally) complex Gaussian signals. We let $z(t) = \mathcal{D}\left[y(t)\right]$ be some nonlinear amplitude distortion and write

$$\mathbb{E}\left[z(t)x^*(t)\right] = g\ \mathbb{E}\left[y(t)x^*(t)\right] \tag{3.1}$$

where $g \in \mathbb{C}$ is some scaling factor.

(a) The general statement. **(b)** Our quantizer context.

Figure 3.1: A visualization of Bussgang's theorem

In order to facilitate the analysis of the effect of one-bit quantization in a communication system, we can frame Bussgang's theorem as in Figure 3.1b. In each timing instant n, the B receive antennas are sampled and the values collected in the vector $\mathbf{y}[n] \in \mathbb{C}^B$. Each entry $y_b[n]$ of the vector $\mathbf{y}[n]$ is then put through the quantizer $\mathcal{Q}$, characterized by (2.4). Collecting all outputs $r_b[n] = \mathcal{Q}(y_b[n])$ in the vector $\mathbf{r}[n] \in \mathbb{C}^B$, we can write (3.1) in a vectorized form as

$$\mathbb{E}\left[\mathbf{r}[n]\mathbf{y}^H[n]\right] = \mathbf{G}\,\mathbb{E}\left[\mathbf{y}[n]\mathbf{y}^H[n]\right] \tag{3.2}$$

where the gain $\mathbf{G}$ is a $B \times B$ diagonal matrix. From (3.2), it follows that $\mathbf{r}[n] = \mathbf{G}\mathbf{y}[n] + \mathbf{d}[n]$, where $\mathbf{d}[n] \in \mathbb{C}^B$, if $\mathbf{d}[n]$ is uncorrelated with $\mathbf{y}[n]$, i.e. $\mathbb{E}\left[\mathbf{y}[n]\mathbf{d}^H[n]\right] = 0$. To verify, we can examine a single entry $r_b[n]$ in $\mathbf{r}[n]$ and substitute on the left-hand side in (3.2). Then,

$$\begin{aligned}
\mathbb{E}\left[([\mathbf{G}]_{b,b}\; y_b[n] + d_b[n])y_b^*[n]\right] &= [\mathbf{G}]_{b,b}\;\mathbb{E}\left[y_b[n]y_b^*[n]\right]\\
\mathbb{E}\left[[\mathbf{G}]_{b,b}\; y_b[n]y_b^*[n] + y_b^*[n]d_b[n]\right] &= [\mathbf{G}]_{b,b}\;\mathbb{E}\left[y_b[n]y_b^*[n]\right]\\
[\mathbf{G}]_{b,b}\;\mathbb{E}\left[y_b[n]y_b^*[n]\right] + \mathbb{E}\left[y_b^*[n]d_b[n]\right] &= [\mathbf{G}]_{b,b}\;\mathbb{E}\left[y_b[n]y_b^*[n]\right]
\end{aligned}$$

where the last step follows as the expectation $\mathbb{E}[\cdot]$ is a linear operator. We see that for the above to hold, the second term on the left-hand side must be equal to zero. Consequently, $d_b[n]$ must be uncorrelated with $y_b[n]$. Consequently, we can indeed write

$$\mathbf{r}[n] = \mathcal{Q}(\mathbf{y}[n]) = \mathbf{G}\mathbf{y}[n] + \mathbf{d}[n]. \tag{3.3}$$

This formulation will become useful when we investigate the effects on the quantized signal in the presence of STO and CFO. The term $\mathbf{d}[n]$ can be viewed as a distortion and captures the adverse effects of the quantizer. Interstingly, while the power of the distortion caused by the one-bit quantizers can be significant in relation to the signal power, illustrated in Figure 1.1d, studies such as [39, 40, 16] has shown that 1) the performance loss is not necessarily as severe as one might intuitively think and 2) only a few bits are required to make the gap within fractions of a dB for low SNRs.

Next, we examine the gain $\mathbf{G}$. A general expression where no particular constraints are placed on the nonlinear amplitude distortion $\mathcal{D}(\cdot)$ can be found in [41], but let us derive the gain in the special case of a one-bit ADC. As mentioned previously $\mathbf{G}$ is a diagonal matrix and we will find the expression for each diagonal entry $[\mathbf{G}]_{b,b}$. For notational clarity, we will here drop the index n and consider a single point in time.

From Equation (3.2), we have

$$[\mathbf{G}]_{b,b} = \frac{\mathbb{E}\left[r_b y_b^*\right]}{\mathbb{E}\left[y_b y_b^*\right]}. \tag{3.4}$$

We can write the numerator of Equation (3.4) as

$$
\begin{aligned}
\mathbb{E}\left[r_b y_b^*\right] &= \mathbb{E}\left[\Big(\mathrm{Re}(r_b) + j\mathrm{Im}(r_b)\Big) \cdot \Big(\mathrm{Re}(y_b) - j\mathrm{Im}(y_b)\Big)\right] \\
&= \mathbb{E}\left[\mathrm{Re}(r_b)\ \mathrm{Re}(y_b)\right] - j\mathbb{E}\left[\mathrm{Re}(r_b)\ \mathrm{Im}(y_b)\right] + j\mathbb{E}\left[\mathrm{Im}(r_b)\ \mathrm{Re}(y_b)\right] + \mathbb{E}\left[\mathrm{Im}(r_b)\ \mathrm{Im}(y_b)\right].
\end{aligned}
$$

Given that y_b is circularly symmetric complex Gaussian variable, the real and imaginary part are uncorrelated, i.e. $\mathbb{E}[\mathrm{Re}(y_b)\ \mathrm{Im}(y_b)] = 0$. From (2.4), note that the real part of r_b has no relation to the imaginary part of y_b, and equivalently for the relation between $\mathrm{Im}(r_b)$ and $\mathrm{Re}(r_b)$. Consequently, as both the real and imaginary parts are zero-mean, we have

$$
\mathbb{E}\left[r_b y_b^*\right] = \mathbb{E}\left[\mathrm{Re}(z_b)\ \mathrm{Re}(y_b)\right] + \mathbb{E}\left[\mathrm{Im}(z_b)\ \mathrm{Im}(y_b)\right].
$$

Since the real and imaginary part of y_b and z_b are identically distributed, we have

$$
\begin{aligned}
\mathbb{E}\left[r_b y_b^*\right] &= 2\ \mathbb{E}\left[\mathrm{Re}(r_b)\ \mathrm{Re}(y_b)\right] \\
&= 2 \int_{y_{b,R}} \frac{1}{\sqrt{2\pi\sigma^2_{y_{b,R}}}} \frac{1}{\sqrt{2}}\ \mathrm{sgn}[y_{b,R}]\ y_{b,R}\ \mathrm{e}^{-y_{b,R}^2/(2\sigma^2_{y_{b,R}})}\ \mathrm{d}y_{b,R}
\end{aligned}
\tag{3.5}
$$

where we have written $\mathrm{Re}(y)$ as $y_{b,R}$ for readability. Denoting $\mathbb{E}\left[y_b y_b^*\right]$ as $\sigma^2_{y_b}$, we can express the distribution of the new variable $y_{b,R}$ as $\mathcal{CN}(0, \sigma^2_{y_b}/2)$. Performing the substitution $\sigma^2_{y_{b,R}} = \sigma^2_{y_b}/2$ and then inserting (3.5) into (3.4) , we find the gain as

$$
\begin{aligned}
[\mathbf{G}]_{b,b} &= \sqrt{\frac{2}{\pi}} \frac{1}{\sigma^2_{y_b}\sqrt{\sigma^2_{y_b}}}\ 2\int_0^\infty y_{b,R}\ \mathrm{e}^{-\frac{y_{b,R}^2}{\sigma^2_{y_b}}}\ \mathrm{d}y_{b,R} \\
&= \sqrt{\frac{2}{\pi}} \frac{1}{\sigma^2_{y_b}\sqrt{\sigma^2_{y_b}}}\ 2\frac{\sigma^2_{y_b}}{2} \\
&= \sqrt{\frac{2}{\pi}} \frac{1}{\sigma_{y_b}}.
\end{aligned}
\tag{3.6}
$$

From (3.6), we see that the gain depends on the second-order statistics of the received signal $\mathbf{y}[n]$. Consequently, the final expression for the gain $\mathbf{G}$ will vary with the assumed channel model and for a frequency-flat channel model and symbol-rate sampling, we have, writing (2.3) in vector notation,

$$
\begin{aligned}
\mathbf{G} &= \sqrt{\frac{2}{\pi}}\left(\operatorname{diag}\left(\mathbb{E}\left[\mathbf{y}[n]\mathbf{y}^H[n]\right]\right)\right)^{-1/2} && (3.7)\\
&= \sqrt{\frac{2}{\pi}}\left(\operatorname{diag}\left(\mathbf{H}\underbrace{\mathbb{E}\left[\mathbf{x}[n]\mathbf{x}^H[n]\right]}_{=\mathbf{I}_U}\mathbf{H}^H + \underbrace{\mathbb{E}\left[\mathbf{w}[n]\mathbf{w}^H[n]\right]}_{=\mathbf{I}_B\cdot N_0}\right)\right)^{-1/2}\\
&= \sqrt{\frac{2}{\pi}}\left(\operatorname{diag}\left(\mathbf{H}\mathbf{H}^H + \mathbf{I}_B\cdot N_0\right)\right)^{-1/2} && (3.8)
\end{aligned}
$$

where $\mathbf{w}[n] \in \mathcal{C}^B$, $\mathbf{I}_U$ and $\mathbf{I}_B$ represents the identity matrix of size U and B, respectively and $\operatorname{diag}(\cdot)$ forms a new diagonal matrix from the main diagonal of a input matrix. The cross-terms in (3.7) are zero as $\mathbf{x}[n]$ are $\mathbf{w}[n]$ are independent and the noise $\mathbf{w}[n]$ is zero-mean.

Lastly, note that the validity of Bussgang's theorem hinges on the fact that the quantizer input, $\mathbf{y}[n] \in \mathbb{C}^B$, is a Gaussian signal. Using the OFDM signalling scheme, this condition can generally be regarded as fulfilled.

3.2 Received signal due to STO and CFO

As mentioned in Section 2.3, the STO will impact the system differently depending on if the DFT window is placed too late or too early. Additionally, the CFO can potentially place significant limitations on the system performance. Our main goal is to find analytic expressions for how the SINDR is affected due to these impairments, as well as the effect of one-bit quantization.

To be able to express the received signal at the bth antenna in a MU-MIMO uplink system with B receive antennas and U UEs with STO and CFO, we begin by defining the set $\mathcal{N} = \{0, 1, \ldots, N-1\}$, i.e. all time-domain samples of an OFDM symbol. Next, we define the set $\mathcal{S}$ to be some subset of $\mathcal{N}$, i.e. $\mathcal{S} \cup \mathcal{S}^c = \mathcal{N}$. Given that the UEs are continuously transmitting OFDM symbols, we use (2.3) and express the ith received symbol at the bth BS antenna as

$$
\begin{aligned}
y_b^{(i)}[n] = \sum_{u=1}^{U} h_{b,u} \mathrm{e}^{j2\pi\epsilon_u n/N} \Big(& x_u^{(i)}[(n+\delta_u)_N]\ \mathbb{1}_{\{\mathcal{S}_u\}}[n] \\
& + x_u^{(i-1)}[n+N+P+\delta_u]\ \mathbb{1}_{\{\mathcal{S}_u^c \wedge \delta_u < -P\}}[n] + x_u^{(i+1)}[n]\ \mathbb{1}_{\{\mathcal{S}_u^c \wedge \delta_u > 0\}}[n]\Big) \\
& + w_b[n] \quad , \quad n = 0, 1, \ldots, N-1 \qquad (3.9)
\end{aligned}
$$

where the $(\cdot)_N$ defines the modulo operator, i.e. $(n)_N = n \bmod N$. Further,

$$\mathcal{S}_u = \begin{cases} \{0, 1, \ldots, N - \delta_u - 1\}, & 0 < \delta_u < N \\ \{d_u, d_u + 1, \ldots, N - 1\}, & -N < \delta_u \leq 0 \end{cases} \tag{3.10}$$

is the set of samples corresponding to the ith transmitted symbol from the uth UE that are received during the ith DFT window, and

$$d_u = \max(-P - \delta_u, 0). \tag{3.11}$$

The parameter d_u captures the number of samples missed if $\delta_u < -P$. As mentioned in Section 2.3, as long as the starting point is taken within the cyclic prefix, no ISI or ICI will be incurred. To keep track of this distinction, we use the parameter d_u. Lastly, as mentioned in 2.3, we will assume that the frame detection preceding the DFT block has limited the STO for the ith symbol to $\pm N$. In fact, as any larger STO than $\pm N$ would mean that no samples of the ith symbol is present in the ith DFT window, it not make much sense pursing an SINDR expression in that case.

As discussed in Section 2.3, the main causes of CFO are a mismatch between the oscillators in the transmitter and receiver, as well as the Doppler effect. We can assume that a single clock is driving all analog components in the receiver, but, naturally, this assumption would not be valid for the transmitting side in a multi-user scenario. Further, we have no reason to assume that all users move at roughly the same speed, meaning that all users have a unique Doppler shift relative to the receiver. Consequently, we must separate the CFO for each user. Similarly, as we have no reason to assume that the users are transmitting in a synchronized fashion equidistantly from the receiver, we also must separate the STO for each user.

As noted in the previous section, the gain $[\mathbf{G}]_{b,b}$ from (3.3) will depend on the second-order statistics of the input, so in order to compute the analytic expression for the received signal, we need to determine $y_b^{(i)}[n]$.

3.2.1 Infinite-precision case

In (3.9), we have used indicator functions to deal with STO larger than or less than zero. We will treat these two cases separately, as they will cause either samples from the next symbol or from the previous one to leak into the current block. In Figure 3.2, the two cases are shown. We will begin with Figure 3.2b.

Case $0 < \delta_u < N$**:**

If the STO is positive, the DFT window will no longer include the complete symbol. We will sample into the cyclic prefix of the next symbol and given a timing offset δ_u, we will miss the $\delta_u - 1$ first samples of the ith symbol and instead get first δ_u symbols of the following OFDM symbol, $i + 1$. Then, the DFT, defined by (2.2), of (3.9) reduces to

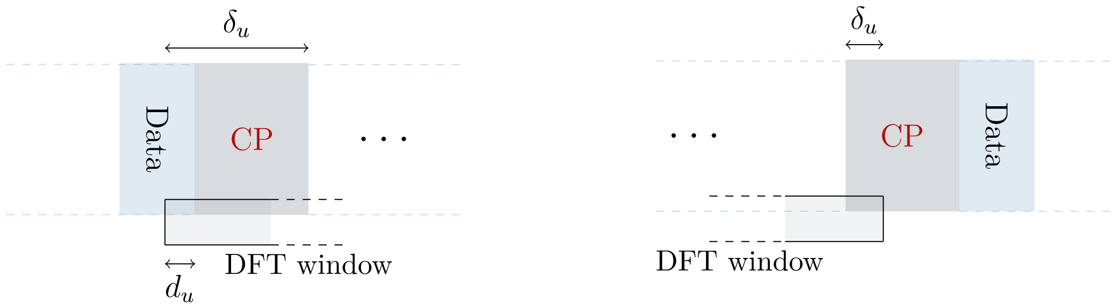

(a) $\delta_u < 0$.

(b) $\delta_u > 0$.

Figure 3.2: Illustration of $\delta_u > 0$ and $\delta_u < 0$.

$$\begin{aligned}\hat{y}_b^{(i)}[k] = \sum_{u=1}^{U}\Bigg(& h_{b,u}\frac{1}{\sqrt{N}}\sum_{n=0}^{N-\delta_u-1} \mathrm{e}^{j2\pi\epsilon_u n/N} x_u^{(i)}\left[(n+\delta_u)_N\right]\ \mathrm{e}^{-j2\pi nk/N} \\ & \underbrace{+\, h_{b,u}\frac{1}{\sqrt{N}}\sum_{n=N-\delta_u}^{N-1} \mathrm{e}^{j2\pi\epsilon_u n/N}\ x_u^{(i+1)}[n]\ \mathrm{e}^{-j2\pi nk/N}\Bigg)}_{=\hat{i}_{b,u}^{\text{ISI, right}}[k]} \\ & + \frac{1}{\sqrt{N}}\sum_{n=0}^{N-1} w_b[n]\ \mathrm{e}^{-j2\pi nk/N}. \end{aligned} \tag{3.12}$$

We note that the last term of (3.12) is the DFT of the thermal noise, $\hat{w}_b[k] = \mathrm{DFT}\{w_b[n]\}$. The term $\hat{i}_{b,u}^{\text{ISI, right}}[k]$ is the ISI caused by the $(i+1)$th symbol during the ith DFT window.

With these definitions, we now have

$$\hat{y}_b^{(i)}[k] = \sum_{u=1}^{U}\left(h_{b,u}\frac{1}{\sqrt{N}}\sum_{n=0}^{N-\delta_u-1} \mathrm{e}^{j2\pi\epsilon_u n/N} x_u^{(i)}\left[(n+\delta_u)_N\right]\ \mathrm{e}^{-j2\pi nk/N} + \hat{i}_{b,u}^{\text{ISI, right}}[k]\right) + \hat{w}_b[k]. \tag{3.13}$$

To continue, we write $x_u^{(i)}[n]$ in (3.13) in the frequency domain using (2.2). Simplifying the exponents and rearranging the sums, we get

$$\hat{y}_b^{(i)}[k] = \sum_{u=1}^{U}\left(h_{b,u}\sum_{k'=0}^{N-1}\hat{x}_u^{(i)}[k']\mathrm{e}^{j2\pi\delta_u k'/N}\ \underbrace{\frac{1}{N}\sum_{n=0}^{N-\delta_u-1}\mathrm{e}^{j2\pi(k'-k+\epsilon_u)n/N}}_{=\psi[k,k']} + \hat{i}_{b,u}^{\text{ISI}}[k]\right) + \hat{w}_b[k]. \tag{3.14}$$

Now, we focus on $\psi[k,k']$. Using the well-known formula for the geometric progression

$$\sum_{n=0}^{N-1} a^n = \frac{1-a^N}{1-a} \tag{3.15}$$

we note that $\psi[k,k']$ bears a strong resemblance to the Euler formula for sine, i.e. $\sin(x) = (\exp(jx) - \exp(-jx))/(2j)$. Multiplying the numerator and denominator with the appropriate values, we get

$$\psi[k,k'] = \mathrm{e}^{j\frac{\pi}{N}(k'-k+\epsilon_u)(N-\delta_u-1)} \; \frac{\sin\left(\frac{\pi}{N}(k'-k+\epsilon_u)(N-\delta_u)\right)}{N\cdot\sin\left(\frac{\pi}{N}(k'-k+\epsilon_u)\right)}. \tag{3.16}$$

Now, consider the case when $k'=k$, i.e. the subcarrier index of interest. Defining

$$g(\alpha,\beta) = \frac{\sin\left(\frac{\pi}{N}(N-\alpha)\beta\right)}{N\cdot\sin(\frac{\pi\beta}{N})} \; \mathrm{e}^{j\frac{\pi}{N}(N-\alpha-1)\beta}, \tag{3.17}$$

note that when $k'=k$, then, from (3.16), $\psi[k,k] = g(\delta_u,\epsilon_u)$. This term captures the attenuation caused by STO, as well as the attenuation and phase shift to due to CFO. Using (3.17), we can now write (3.14) as

$$\begin{aligned}\hat{y}_b^{(i)}[k] = \sum_{u=1}^{U}\Bigg(& h_{b,u}\; g(\delta_u,\epsilon_u)\; \mathrm{e}^{j2\pi\delta_u k/N}\; \hat{x}_u^{(i)}[k] \\ & + \underbrace{h_{b,u} \sum_{\substack{k'=0\\k'\neq k}}^{N-1} \frac{\sin\left(\frac{\pi}{N}(N-\delta_u)(k'-k+\epsilon_u)\right)}{N\cdot\sin(\frac{\pi}{N}(k'-k+\epsilon_u))}\; \mathrm{e}^{j\frac{\pi}{N}(N-\delta_u-1)(k'-k+\epsilon_u)}\; \mathrm{e}^{j2\pi\delta_u k'/N}\hat{x}_u^{(i)}[k']}_{=\phi[k]} \\ & + \hat{i}_{b,u}^{\text{ISI, right}}[k]\Bigg) + \hat{w}_b[k]. \end{aligned} \tag{3.18}$$

Now, we perform the variable change $k'' = k - k'$ and write $\phi[k]$ as

$$\phi[k] = \sum_{k''=1}^{N-1} \frac{\sin\left(\frac{\pi}{N}(N-\delta_u)(\epsilon_u-k'')\right)}{N\cdot\sin\left(\frac{\pi}{N}(\epsilon_u-k'')\right)}\; \mathrm{e}^{j\frac{\pi}{N}(N-\delta_u-1)(\epsilon_u-k'')}\; \mathrm{e}^{j2\frac{\pi}{N}(k-k'')\delta_u}\hat{x}_u^{(i)}\left[(k-k'')_N\right]. \tag{3.19}$$

The first two factors of (3.19) is precisely $g(\delta_u, \epsilon_u - k'')$, so finally, (3.18) becomes

$$\begin{aligned}\hat{y}_b^{(i)}[k] = \sum_{u=1}^{U} \Big(& h_{b,u}\ g(\delta_u, \epsilon_u)\ \mathrm{e}^{j2\pi\delta_u k/N}\ \hat{x}_u^{(i)}[k] \\ & \underbrace{+ \sum_{k''=1}^{N-1} h_{b,u}\ g(\delta_u, \epsilon_u - k'')\ \mathrm{e}^{j2\pi\delta_u (k-k'')/N}\ \hat{x}_u^{(i)}[(k-k'')_N]}_{=\hat{i}_{b,u}^{\text{ICI, right}}[k]} \\ & + \hat{i}_{b,u}^{\text{ISI, right}}[k] \Big) + \hat{w}_b[k]. \end{aligned} \tag{3.20}$$

Looking at the terms of (3.20), we see the expected ICI terms due to loss of subcarrier orthogonality as well as ISI due to the next symbol.

Case $-P < \delta_u < 0$**:**

Looking at Figure 3.2a, we see that in a flat-fading channel model, there will be no delay spread caused by the preceding symbol that will distort the cyclic prefix of the current symbol. This means that as long as $-P \leq \delta_u \leq 0$, there will be no interference from the previous symbol. As mentioned in Section 2.3, the only effect will be a rotation of the received constellation, as the orthogonality between the subcarriers is not affected. If δ_u is in this interval, then d_u in (3.11) will be zero and from (3.9), we have

$$y_b^{(i)}[n] = \sum_{u=1}^{U} h_{b,u}\ \mathrm{e}^{j2\pi\epsilon_u n/N}\ x_u^{(i)}\left[(n+\delta_u)_N\right] + w_b[n]. \tag{3.21}$$

Applying a DFT to (3.21) and proceeding in a similar manner as in the previous case, we arrive at

$$\hat{y}_b^{(i)}[k] = \sum_{u=1}^{U} h_{b,u}\ g(0, \epsilon_u)\ \mathrm{e}^{j2\pi\delta_u k/N}\ \hat{x}_u^{(i)}[k] + \hat{w}_b[k]. \tag{3.22}$$

We see that the only effect from the STO is a rotation of the received symbol.

Case $-N < \delta_u < -P$**:**

Here, we will miss the first d_u samples of symbol i and instead get the last d_u samples of symbol $i-1$. This means that we will get ISI. Applying the DFT to (3.9), we get

$$\hat{y}_b^{(i)}[k] = \sum_{u=1}^{U} \Bigg(h_{b,u} \frac{1}{\sqrt{N}} \sum_{n=d_u}^{N-1} \mathrm{e}^{j2\pi\epsilon_u n/N} x_u^{(i)} \left[(n+\delta_u)_N\right] \ \mathrm{e}^{-j2\pi nk/N}$$
$$+ \underbrace{h_{b,u} \frac{1}{\sqrt{N}} \sum_{n=0}^{d_u-1} \ \mathrm{e}^{j2\pi\epsilon_u n/N} \ x_u^{(i-1)}[n+N-d_u] \ \mathrm{e}^{-j2\pi nk/N} \Bigg)}_{=\hat{i}_{b,u}^{\text{ISI, left}}[k]}$$
$$+ \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} \ w_b[n] \ \mathrm{e}^{-j2\pi nk/N}. \tag{3.23}$$

Using the same steps in the case of $0 < \delta_u < N$, we arrive at

$$\hat{y}_b^{(i)}[k] = \sum_{u=1}^{U} \left(h_{b,u} \sum_{k'=0}^{N-1} \hat{x}_u^{(i)}[k'] \mathrm{e}^{j2\pi\delta_u k'/N} \ \underbrace{\frac{1}{N} \sum_{n=d_u}^{N-1} \ \mathrm{e}^{j2\pi(k'-k+\epsilon_u)n/N}}_{=\xi[k,k']} + \hat{i}_{b,u}^{\text{ISI, left}}[k] \right) + \hat{w}_b[k]. \tag{3.24}$$

Again, we focus at the innermost sum over n and write

$$\begin{aligned}
\xi[k,k'] &= \frac{1}{N} \sum_{n=d_u}^{N-1} \ \mathrm{e}^{j2\pi(k'-k+\epsilon_u)n/N} \\
&= \frac{1}{N} \left(\sum_{n=0}^{N-1} \mathrm{e}^{j2\pi(k'-k+\epsilon_u)n/N} - \sum_{n=0}^{d_u-1} \mathrm{e}^{j2\pi(k'-k+\epsilon_u)n/N} \right) \\
&= \frac{1}{N} \left(\frac{1-\mathrm{e}^{j2\pi(k'-k+\epsilon_u)}}{1-\mathrm{e}^{j2\pi(k'-k+\epsilon_u)/N}} - \frac{1-\mathrm{e}^{j2\pi(k'-k+\epsilon_u)d_u/N}}{1-\mathrm{e}^{j2\pi(k'-k+\epsilon_u)/N}} \right) \\
&= \frac{1}{N} \cdot \frac{\mathrm{e}^{j2\pi(k'-k+\epsilon_u)d_u/N} - \mathrm{e}^{j2\pi(k'-k+\epsilon_u)}}{1-\mathrm{e}^{2\pi(k'-k+\epsilon_u)/N}}.
\end{aligned} \tag{3.25}$$

Multiplying (3.25) with the appropriate factors, we can again use Euler's formula for sine to find

$$\xi[k,k'] = \frac{\sin\left(\frac{\pi}{N}(k'-k+\epsilon_u)(N-d_u)\right)}{N \cdot \sin\left(\frac{\pi}{N}(k'-k+\epsilon_u)\right)} \ \mathrm{e}^{j\pi(k'-k+\epsilon_u)(N-d_u-1)/N} \ \mathrm{e}^{-j2\pi(k'-k+\epsilon_u)d_u/N}. \tag{3.26}$$

For the case $k' = k$, (3.26) simplifies to

$$\xi[k,k] = g(d_u,\epsilon_u)\mathrm{e}^{-j2\pi\epsilon_u d_u/N}. \tag{3.27}$$

Using the same change of variable $k'' = k - k'$ as before, we can finally write (3.23) as

$$\begin{aligned}\hat{y}_b^{(i)}[k] = \sum_{u=1}^{U} \Big(& h_{b,u}\ g(d_u, \epsilon_u)\ \mathrm{e}^{j2\pi\delta_u k/N}\ \mathrm{e}^{-j2\pi\epsilon_u\delta_u/N}\ \hat{x}_u^{(i)}[k] \\ & \underbrace{+ \sum_{k''=1}^{N-1} h_{b,u}\ g(d_u, \epsilon_u - k'')\ \mathrm{e}^{j2\pi\delta_u(k-k'')/N}\ \mathrm{e}^{-j2\pi(k''+\epsilon_u)d_u/N}\ \hat{x}_u^{(i)}[(k-k'')_N]}_{=\hat{i}_{b,u}^{\text{ICI, left}}[k]} \\ & + \hat{i}_{b,u}^{\text{ISI, left}}[k] \Big) + \hat{w}_b[k]. \end{aligned} \tag{3.28}$$

Comparing Equation (3.20) and (3.28), we note that they are highly similar. Again, we get ISI due to the misaligned window, as well as ICI since the subcarrier orthogonality is not preserved. This is to be expected, as the effect in either direction largely amounts to the same thing.

To verify our calculations, we examine the expression when either δ_u or ϵ_u is set to zero. Beginning with no CFO, we get

$$g(\delta_u, 0) = \frac{N - \delta_u}{N}$$

using l'Hospital's rule. For $g(\delta_u, -k'')$, we get

$$\begin{aligned} g(\delta_u, -k'') &= \frac{1}{N} \cdot \frac{1 - \mathrm{e}^{-j2\pi k''}\mathrm{e}^{j2\pi\delta_u k''/N}}{\mathrm{e}^{j\pi k''/N} - \mathrm{e}^{-j\pi k''/N}} \mathrm{e}^{j\pi k''/N} \\ &= \frac{1}{N} \cdot \frac{1 - \mathrm{e}^{j2\pi\delta_u k''/N}}{1 - \mathrm{e}^{j2\pi k''/N}} \end{aligned}$$

where we used that $\mathrm{e}^{j2\pi k''} = 1$, since $k'' \in \mathbb{Z}$.

With these results, (3.20) reduces to

$$\begin{aligned}\hat{y}_b^{(i)}[k] = \sum_{u=1}^{U} \Big(& \frac{N-\delta_u}{N} h_{b,u} \mathrm{e}^{j2\pi\delta_u k/N}\ h_{b,u}\ \hat{x}_u^{(i)}[k] \\ & + \frac{1}{N} \sum_{k''=1}^{N-1} h_{b,u} \frac{1 - \mathrm{e}^{j2\pi k''\delta_u/N}}{1 - \mathrm{e}^{j2\pi k''/N}}\ \mathrm{e}^{j2\pi\delta_u(k-k'')/N}\ \hat{x}_u^{(i)}[(k-k'')_N] \\ & + \hat{i}_{b,u}^{\text{ISI, right}}[k] \Big) + \hat{w}_b[k]. \end{aligned} \tag{3.29}$$

This is the same result stated in for example [42, cf. (1)]. Similarly, (3.28) reduces to [42, cf. (3)] if $\epsilon_u = 0$.

Setting δ_u (or d_u) to zero, we get,

$$g(0, \epsilon_u) = \frac{\sin(\pi\epsilon_u)}{N \cdot \sin(\frac{\pi\epsilon_u}{N})} \, \mathrm{e}^{j\pi\epsilon_u(N-1)/N}$$

Using this, we can reduce (3.20) and (3.28) to

$$\begin{aligned}\hat{y}_b^{(i)}[k] = \sum_{u=1}^{U} \Bigg(& \frac{\sin(\pi\epsilon_u)}{N \cdot \sin(\frac{\pi\epsilon_u}{N})} \, \mathrm{e}^{j\pi\epsilon_u(N-1)/N} \, \hat{x}_u^{(i)}[k] \\ & + \sum_{k''=1}^{N-1} h_{b,u} \, \frac{\sin(\pi(\epsilon_u - k''))}{N \cdot \sin(\frac{\pi}{N}(\epsilon_u - k''))} \, \mathrm{e}^{j\pi(\epsilon_u - k'')(N-1)/N} \hat{x}_u^{(i)} \left[(k - k'')_N\right] \Bigg) \\ & + \hat{w}_b[k]. \end{aligned} \tag{3.30}$$

The expression (3.30) can be compared to for example [32], where a similar expression was found for the case $h_{b,u} = 1$.

3.2.2 One-bit quantization

Now that we have found the expressions for $y_b[n]$ and $\hat{y}_b[k]$, we are ready to continue with the one-bit quantizer.

From (3.6), we see that we need to determine the power of the received signal on each antenna, $\sigma^2_{y_b}$. Equivalently, we can examine the power of $\hat{y}_b[k] = \mathrm{DFT}\{y_b[n]\}$, given the DFT definition in (2.2). Moreover, in this particular setting, it is trivial to see that

$$\begin{aligned}\hat{r}_b^{(i)}[k] &= \mathrm{DFT}\{r_b^{(i)}[n]\} \\ &= [\mathbf{G}]_{b,b}\hat{y}_b^{(i)}[k] + \hat{d}_b[k]\end{aligned} \tag{3.31}$$

where $\hat{d}_b[k] = \mathrm{DFT}\{d_b[n]\}$ and $[\mathbf{G}]_{b,b}$ is given by (3.6). From (3.20) and (3.28), we see that the signal $\hat{y}_b^{(i)}[k]$ on antenna b after the DFT has the following form

$$\hat{y}_b^{(i)}[k] = \hat{s}_{b,u}[k] + \hat{i}^{\mathrm{ISI}}_{b,u}[k] + \hat{i}^{\mathrm{ICI}}_{b,u}[k] + \underbrace{\sum_{\substack{u'=1 \\ u' \neq u}}^{U} \hat{s}_{b,u'}[k] + \hat{i}^{\mathrm{ISI}}_{b,u'}[k] + \hat{i}^{\mathrm{ICI}}_{b,u'}[k]}_{\hat{i}^{\mathrm{MUI}}_{b,u}[k]} + \hat{w}_b[k] \tag{3.32}$$

where

$$\hat{i}^{\text{ISI}}_{b,u}[k] = \begin{cases} \hat{i}^{\text{ISI, right}}_{b,u}[k], & 0 < \delta_u < N \\ \hat{i}^{\text{ISI, left}}_{b,u}[k], & -N < \delta_u < -P \end{cases} \tag{3.33}$$

and

$$\hat{i}^{\text{ICI}}_{b,u}[k] = \begin{cases} \hat{i}^{\text{ICI, right}}_{b,u}[k], & 0 < \delta_u < N \\ \hat{i}^{\text{ICI, left}}_{b,u}[k], & -N < \delta_u < -P. \end{cases} \tag{3.34}$$

Furthermore, $\hat{s}_{b,u}[k]$ defines the part of the received signal from the uth UE, i.e.

$$\hat{s}_{b,u}[k] = h_{b,u}\ \gamma(\delta_u, \epsilon_u)\ \hat{x}^{(i)}_u[k]\ \Psi[\delta_u] \tag{3.35}$$

where

$$\gamma(\delta_u, \epsilon_u) = \begin{cases} g(\delta_u, \epsilon_u), & 0 < \delta_u < N \\ g(d_u, \epsilon_u), & -N \leq \delta_u \leq 0 \end{cases} \tag{3.36}$$

and

$$\Psi[\delta_u] = \begin{cases} \mathrm{e}^{j2\pi\delta_u k/N}, & -P \leq \delta_u < N \\ \mathrm{e}^{j2\pi\delta_u k/N}\ \mathrm{e}^{-j2\pi\epsilon_u\delta_u/N}, & -N < \delta_u < -P. \end{cases} \tag{3.37}$$

In (3.32), $\hat{i}^{\text{ISI}}_{b,u}[k]$ and $\hat{i}^{\text{ICI}}_{b,u}[k]$ represents the interference, stemming from the uth UE on the bth antenna, caused by STO and CFO. In addition to these interferences, there is also MU interference, denoted with $\hat{i}^{\text{MUI}}_{b,u}[k]$, as well as thermal noise, $\hat{w}_b[k]$.

3.2.3 ZF equalization

Inserting (3.32) into (3.31), and inserting (3.31) into (2.5), we get that the estimated frequency-domain symbol on from the uth UE during the ith DFT window, which can be written as

$$\begin{aligned} \hat{x}^{\text{est},(i)}_u[k] &= \mathbf{a}_u^T \mathbf{G}\mathbf{h}_u \gamma(\delta_u, \epsilon_u)\ \hat{x}^{(i)}_u[k]\Psi[\delta_u] \\ &\quad + \mathbf{a}_u^T \mathbf{G}\hat{\mathbf{i}}^{\text{ISI}}_u[k] + \mathbf{a}_u^T \mathbf{G}\hat{\mathbf{i}}^{\text{ICI}}_u[k] + \mathbf{a}_u^T \mathbf{G}\hat{\mathbf{i}}^{\text{MUI}}_u[k] \\ &\quad + \mathbf{a}_u^T \hat{\mathbf{d}}[k] + \mathbf{a}_u^T \mathbf{G}\hat{\mathbf{w}}[k] \end{aligned} \tag{3.38}$$

where $\mathbf{a}_u = [a_{u,1}, a_{u,2}, \ldots, a_{u,B}]^T$ is the uth row of the ZF equalization matrix $\mathbf{A}$ defined in (2.6) and $\mathbf{h}_u = [h_{1,u}, h_{2,u}, \ldots, h_{B,u}]^T$ is the uth column of the channel matrix $\mathbf{H}$. We have further defined $\hat{\mathbf{i}}_u^{\text{ISI}}[k] = \left[\hat{i}_{1,u}^{\text{ISI}}[k], \hat{i}_{2,u}^{\text{ISI}}[k], \ldots, \hat{i}_{B,u}^{\text{ISI}}[k]\right]^T$, $\hat{\mathbf{i}}_u^{\text{ICI}}[k] = \left[\hat{i}_{1,u}^{\text{ICI}}[k], \hat{i}_{2,u}^{\text{ICI}}[k], \ldots, \hat{i}_{B,u}^{\text{ICI}}[k]\right]^T$, $\hat{\mathbf{i}}_u^{\text{MUI}}[k] = \left[\hat{i}_{1,u}^{\text{MUI}}[k], \hat{i}_{2,u}^{\text{MUI}}[k], \ldots, \hat{i}_{B,u}^{\text{MUI}}[k]\right]^T$, $\hat{\mathbf{d}}[k] = \left[\hat{d}_1[k], \hat{d}_2[k], \ldots, \hat{d}_B[k]\right]^T$, and $\hat{\mathbf{w}}[k] = [\hat{w}_1[k], \hat{w}_2[k], \ldots, \hat{w}_B[k]]^T$.

From (3.38), we find that the SINDR at the uth UE can be written as

$$\text{SINDR}_u = \frac{|\gamma(\delta_u, \epsilon_u)|^2 |\mathbf{a}_u^T \mathbf{G} \mathbf{h}_u|^2}{I_u^{\text{ISI}} + I_u^{\text{ICI}} + I_u^{\text{MUI}} + \mathbf{a}_u^T \mathbb{E}\left[\hat{\mathbf{d}}[k]\hat{\mathbf{d}}^H[k]\right] \mathbf{a}_u^* + N_0 \mathbf{a}_u^T \mathbf{G}\mathbf{G}^H \mathbf{a}_u^*} \tag{3.39}$$

since the interference sources, namely the ISI, ICI, and MU interference, quantization distortion, and thermal noise are uncorrelated. The numerator in (3.39) is the power of the received signal from the uth UE, i.e. the first term in (3.38). For the terms in the denominator, we will now examine them individually.

Power of I_u^{ISI}:

To compute the power in the ISI term, we need to distinguish between two cases. If $0 < \delta < N$, the DFT window includes δ_u samples of the next symbol. If $-N < \delta_u < -P$, d_u, the DFT window includes d_u samples of the previous symbol. In either case, the power of this term will be a fraction of the power of a full symbol, directly tied to the STO. The power of a full symbol simply $|\mathbf{a}_u^T \mathbf{G} \mathbf{h}_u|^2$, so we find

$$I_u^{\text{ISI}} = \mathbb{E}\left[|\hat{\mathbf{i}}_u^{\text{ISI}}[k]|^2\right] = \begin{cases} \frac{\delta_u}{N} \, |\mathbf{a}_u^T \mathbf{G} \mathbf{h}_u|^2, & \delta_u > 0 \\ \frac{d_u}{N} \, |\mathbf{a}_u^T \mathbf{G} \mathbf{h}_u|^2, & \delta_u \leq 0. \end{cases} \tag{3.40}$$

Power of I_u^{ICI}:

Starting from (3.20),

$$\begin{aligned}
\mathbb{E}\left[|\mathbf{a}_u^T \mathbf{G} \hat{\mathbf{i}}_u^{\text{ICI, right}}[k]|^2\right] &= |\mathbf{a}_u^T \mathbf{G} \mathbf{h}_u|^2 \sum_{k'=1}^{N-1} \sum_{k''=1}^{N-1} g(\delta_u, \epsilon_u - k') \; g^*(\delta_u, \epsilon_u - k'') \\
&\quad \cdot \underbrace{\mathbb{E}\left[\hat{x}_u^{(i)}[(k-k')_N] \hat{x}_u^{*,(i)}[(k-k'')_N]\right]}_{= \begin{cases} 1, & k' = k'' \\ 0, & k' \neq k'' \end{cases}, \text{ since symbol-sampling}} \\
&= |\mathbf{a}_u^T \mathbf{G} \mathbf{h}_u|^2 \sum_{k'=1}^{N-1} \left[\frac{\sin(\frac{\pi}{N}(N-\delta_u)(\epsilon_u - k'))}{N \cdot \sin(\frac{\pi}{N}(\epsilon_u - k'))}\right]^2 \\
&= |\mathbf{a}_u^T \mathbf{G} \mathbf{h}_u|^2 \left(\sum_{k'=0}^{N-1} \left[\frac{\sin(\frac{\pi}{N}(N-\delta_u)(\epsilon_u - k'))}{N \cdot \sin(\frac{\pi}{N}(\epsilon_u - k'))}\right]^2 - \left[\frac{\sin(\frac{\pi}{N}(N-\delta_u)\epsilon_u)}{N \cdot \sin(\frac{\pi \epsilon_u}{N})}\right]^2 \right).
\end{aligned} \tag{3.41}$$

In (3.41), note that the second term is precisely $g|(\delta_u, \epsilon_u)|^2$. For the first term, we find

$$\sum_{k'=0}^{N-1} \left[\frac{\sin(\frac{\pi}{N}(N-\delta_u)(\epsilon_u - k'))}{N \cdot \sin(\frac{\pi}{N}(\epsilon_u - k'))} \right]^2 = \frac{N-\delta_u}{N}. \tag{3.42}$$

Hence,

$$\mathbb{E}\left[|\mathbf{a}_u^T \mathbf{G} \hat{\mathbf{i}}_u^{\text{ICI, right}}[k]|^2\right] = |\mathbf{a}_u^T \mathbf{G} \mathbf{h}_u|^2 \left[\frac{N-\delta_u}{N} - |g(\delta_u, \epsilon_u)|^2 \right]. \tag{3.43}$$

Starting from (3.28) instead, we obtain

$$\mathbb{E}\left[|\mathbf{a}_u^T \mathbf{G} \hat{\mathbf{i}}_u^{\text{ICI, left}}[k]|^2\right] = |\mathbf{a}_u^T \mathbf{G} \mathbf{h}_u|^2 \left[\frac{N-d_u}{N} - |g(d_u, \epsilon_u)|^2 \right]. \tag{3.44}$$

Examining (3.43) and (3.44), we note that we can use $\gamma(\delta_u, \epsilon_u)$ defined in (3.36) to write

$$I_u^{\text{ICI}} = \mathbb{E}\left[|\mathbf{a}_u^T \mathbf{G} \hat{\mathbf{i}}_u^{\text{ICI}}[k]|^2\right] = |\mathbf{a}_u^T \mathbf{G} \mathbf{h}_u|^2 \left[\gamma(\delta_u, 0) - |\gamma(\delta_u, \epsilon_u)|^2 \right]. \tag{3.45}$$

Power of I_u^{MUI}:

As the transmitted signal is normalized, the multi-user interference will only depend on the power in the channel, i.e.

$$I_u^{\text{MUI}} = \mathbb{E}\left[|\mathbf{a}_u^T \mathbf{G} \hat{\mathbf{i}}_u^{\text{MUI}}|^2\right] = \sum_{\substack{u'=1 \\ u' \neq u}}^{U} |\mathbf{a}_u^T \mathbf{G} \mathbf{h}_u|^2. \tag{3.46}$$

Power of quantization distortion:

To find the power of the quantization distortion, we look at (3.3). This describes the quantizer input-output relationship. Rearranging, we find

$$\mathbb{E}\left[\mathbf{d}[k]\mathbf{d}^H[k]\right] = \underbrace{\mathbb{E}\left[\mathbf{r}[k]\mathbf{r}^H[k]\right]}_{=\mathbf{C_{rr}}} - \mathbf{G}\, \underbrace{\mathbb{E}\left[\mathbf{y}[k]\mathbf{y}^H[k]\right]}_{=\mathbf{C_{yy}}} \mathbf{G} \tag{3.47}$$

where $\mathbf{C_{rr}}$ and $\mathbf{C_{yy}}$ are used to denote the covariance matrix of the output and inputs to the quantizer, respectively. The covariance of the input in time-domain was given in (3.8). Computing the covariance matrix $\mathbf{C_{rr}}$ is in the special case of one-bit quantizers fairly straightforward, as we can resort to the arcsine law [43]. Using this law, we find

$$\mathbf{C_{rr}} = \frac{2}{\pi} \left(\arcsin\left(\mathbf{D_y}^{-1/2}\mathrm{Re}(\mathbf{C_{yy}})\mathbf{D_y}^{-1/2}\right) + j \arcsin\left(\mathbf{D_y}^{-1/2}\mathrm{Im}(\mathbf{C_{yy}})\mathbf{D_y}^{-1/2}\right)\right) \tag{3.48}$$

where $\mathbf{D_y}$ refers to the diagonal elements of $\mathbf{C_{yy}}$, i.e. $\mathbf{D_y} = \mathrm{diag}\left(\mathbf{C_{yy}}\right)$. Consequently, we can write the power of the distortion term as

$$\mathbb{E}\left[\mathbf{d}[k]\mathbf{d}^H[k]\right] = \mathbf{C_{rr}} - \mathbf{G}\mathbf{C_{yy}}\mathbf{G}. \tag{3.49}$$

In (3.39), we have the quantization distortion in the frequency-domain. However, due to Parseval's theorem, we can examine the power in either the time- or frequency-domain if the DFT operation is defined as in (2.2). Hence, we can use (3.49) to express the power of the quantization distortion.

Lastly, the power of the noise is found as

$$\begin{aligned}\mathbb{E}\left[(\mathbf{a}_u^T\mathbf{G}\hat{\mathbf{w}}[k])(\mathbf{a}_u^T\mathbf{G}\hat{\mathbf{w}}[k])^H\right] &= \mathbf{a}_u^T\mathbf{G}\mathbb{E}\left[\hat{\mathbf{w}}[k]\hat{\mathbf{w}}[k]^H\right]\mathbf{G}\mathbf{a}_u^* \\ &= N_0\mathbf{a}_u^T\mathbf{G}\mathbf{G}^H\mathbf{a}_u^* \end{aligned} \tag{3.50}$$

where the second equality follows as the noise is white and zero mean.

Inserting the computed power terms (3.40), (3.45) and (3.46) into (3.39), we find a closed-form expression for the SINDR for a given STO and CFO.

4

Synchronization in OFDM systems

From the analysis in Section 3, it is clear that the effects of both STO and CFO can potentially be quite severe in the context of OFDM. The inherent sensitivity to these errors is something that has been known for a long time. For example in [44], it was demonstrated that OFDM is several orders of magnitude more sensitive to frequency offsets than single-carrier modulation. This is not particularly surprising, as orthogonality between the subchannels is instantly destroyed if the received signal is demodulated with the wrong frequencies. Consequently, it is critical to devise strategies to estimate these offsets as accurately as possible.

Early papers resembling the methods of today in this field, such as [45, 46, 47], demonstrated different approaches to both temporally locate the signal, as well as estimating its carrier frequency. A ground-breaking paper in the area of synchronization for OFDM was presented by in 1997 Schmidl and Cox in [48], even if the fundamental ideas are presented already in 1996 [49]. They expanded on the method in [46] and their work found its way into several standards, such as the wireless local area network (WLAN) standard 802.11a [50]. This method will be examined in detail and an extension developed in [51], as well as a method relying on the cyclic prefix presented in [52, 53] and, lastly, a more modern method reminiscent of how synchronization is achieved in the current generation of mobile systems. Before delving into the different methods of synchronization, we will take a brief look at the history of synchronization itself.

4.1 Background

At the most fundamental level, timing synchronization requires the localization of a specific sequence within a window of samples. For example, in a system transmitting raw bits, you could devise a system where every payload is preceded by a known sequence. If we could accurately locate the position of the synchronization sequence, we would also have located the payload. Even though a system transmitting raw bits might seem far from the OFDM system considered in this thesis, it is a good starting point in order to say something general about synchronization.

The natural questions to ask in this setting would be how to go about finding a known sequence within a window and, next, whether the design of the synchronization sequence has any influence on our ability to locate it. Going further, it could also be questioned whether the synchronization sequence needs to be explicitly known at the receiver.

Regarding how to determine the frequency offset of a signal, it hinges on accurate estimates of the temporal position of a sequence. Say for example that a number of known symbols are transmitted and then correctly located on the receiving side. By comparing the phase differences of the symbols in the known sequence on the transmitting and receiving sides, we are able to extract information about the frequency offset. As the timing estimation is key to extract information of a possible frequency offset, the following discussion will focus on the timing aspect.

Optimally locating a sequence

To the first question, the standard method was (and is) to find the peak in the correlation spectrum produced when correlating the window of samples with the known sequence. Whether or not this is optimal and how the nature of the samples in which the sequence is embedded might have an influence, was not definitely settled until [54]. The set-up considered is depicted in Figure 4.1a, where the top part shows the signal $\mathbf{x} \in \mathcal{X}^N, \mathcal{X} = \{-1, 1\}$ to be transmitted over an AWGN channel.

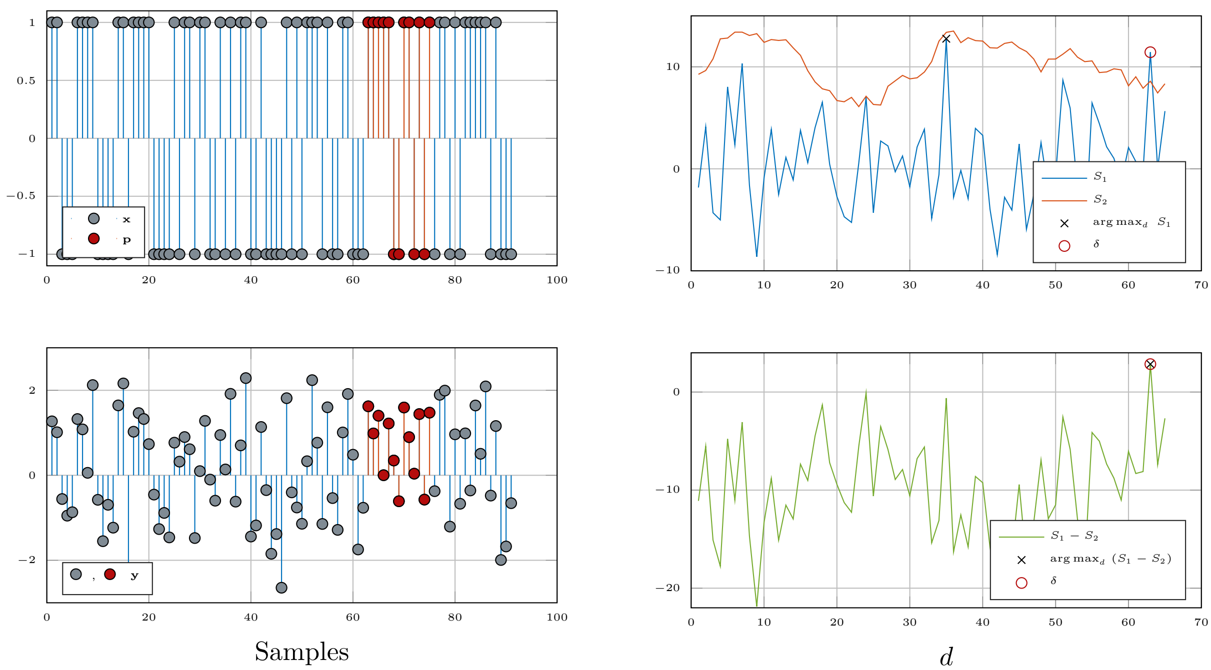

(a) The context in [54], SNR = 0 dB.

(b) The output of the correlation and the optimum rule.

Figure 4.1: The optimum rule for detecting a known sequence under the influence of AWGN.

In $\mathbf{x}$, a known sequence $\mathbf{p}$ of length $L < N$ is embedded, beginning at an unknown location δ. The sequence $\mathbf{x}$ is then subjected to noise, forming $\mathbf{y} = \mathbf{x} + \mathbf{w}$, where $\mathbf{w} \sim \mathcal{N}(0, N_0/2)$. This is depicted in the bottom part of 4.1a and the task is to precisely locate known sequence. Starting from these assumption, Massey derived the optimal rule for the estimation of the position δ^{est} as

$$\delta^{\text{est}} = \arg\max_{d} \left[\underbrace{\sum_{m=0}^{L-1} p[m]y[m+d]}_{S_1[d]} - \underbrace{\sum_{m=0}^{L-1} f(y[m+d])}_{S_2[d]} \right] \tag{4.1}$$

We see that the first term S_1 of (4.1) is indeed a correlation term, in line with the expectations at the time. In addition to the correlation term, there is also a second term S_2 present. This a correction term that accounts for the power in the received signal $\mathbf{y}$. With the assumption that a sequence of ± 1s given by $\mathbf{x}$ is transmitted, S_2 is depicted in the top part of Figure 4.1b, together with S_1. We see that if we only maximize S_1, which will be referred to as the correlation rule, it will produce an erroneous decision regarding the true position of the synchronization sequence, as the peak at $\delta^{\text{est}} = 35$ is greater than the true position $\delta = 63$. However, when adding the correction term, as in the bottom part of Figure 4.1b, the correct location is indeed found.

In Figure 4.2a, the probability of $\delta^{\text{est}} \neq \delta$ is depicted. We see that in the SNR region near 0 dB, Massey's optimum rule defined by (4.1) provides an approximate 3 dB gain over the correlation rule and around twice that a few dBs higher. Figure 4.2a also depicts a high SNR approximation for the correction term, which has a very similar performance as the optimum rule. In this setting, the optimal S_2^{opt} has the form

$$S_2^{\text{opt}}[d] = \frac{N_0}{2\sqrt{2}} \sum_{m=0}^{L-1} \ln\left(\cosh\left(\sqrt{2}/N_0\; y[m+d]\right)\right) \tag{4.2}$$

The computational complexity of (4.2) in apparent, involving both the $\ln(\cdot)$ and $\cosh(\cdot)$ functions. The high SNR approximation, on the other hand, takes the form

$$S_2^{\text{approx.}}[d] = \sum_{m=0}^{L-1} |y[m+d]| \tag{4.3}$$

Consequently, it requires only $L(N-L)$ additions.

It is quite fascinating that the high SNR approximation of the correction term seems to yield comparable performance to when using the optimum S_2^{opt}. This was noted already in Massey's original paper, and conclusively confirmed after extensive simulations in [55]. Going further on the topic of simplifications of the optimal rule, it has been argued in for example [56], that as the received signal in a real system is always constrained to some fixed interval, owing to the AGC/ADC, the correction term will vary very little. Hence, it can in most practical cases be disregarded. As we will see in for example Section 4.3, a variant of the correlation rule without calculating any correction term is used to find the STO.

Another interesting fact regarding the optimum rule is that it becomes the correlation rule in the case of one-bit quantization, depicted in Figure 4.2b. Here, we form $\mathbf{r} = \mathcal{Q}(\mathbf{y}) = \pm 1/\sqrt{(2)}$. In

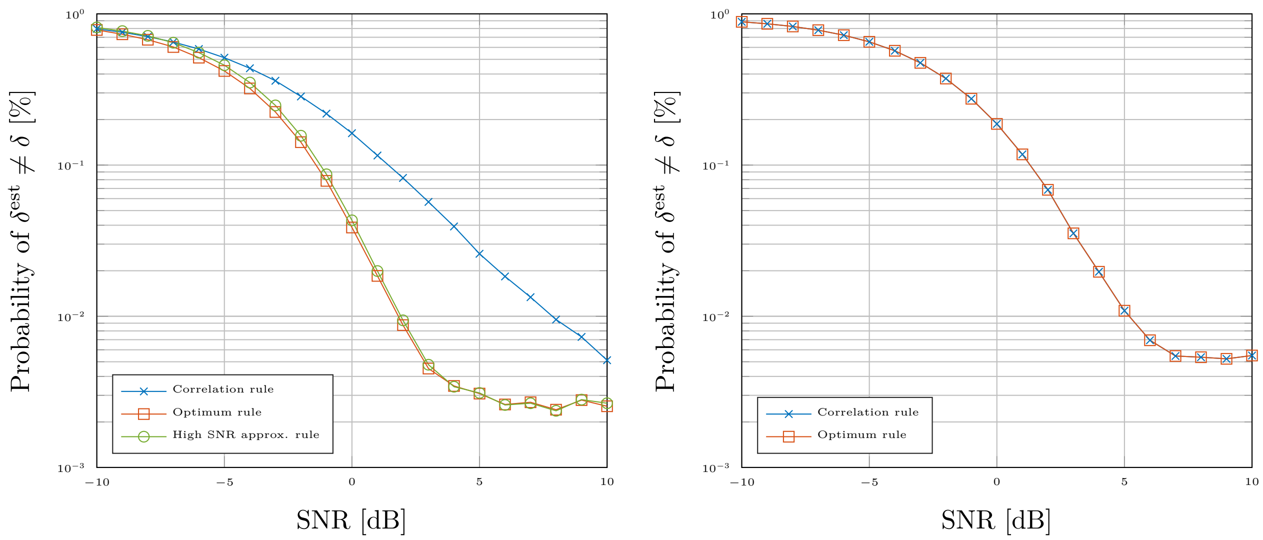

(a) The probability of $\delta^{\text{est}} \neq \delta$.

(b) The error probability with one-bit quantization.

Figure 4.2: The probabilities of an erroneous decision regarding the true position δ of the synchronization sequence. Results obtained via simulation the detection of a length 13 Barker sequence over the AWGN channel 10^5 times.

this particular setting, it can be understood from the fact that as $\cosh(\cdot)$ from (4.2) is an even function and as $\mathrm{r} = \pm 1$, S_2^{opt} will be the same for all values of d. An explanation on a more intuitive level is to consider the purpose of the correction term. It is intended to take the data surrounding the synchronization sequence into account weaken peaks in the correlation spectrum due to particularly powerful input. In the case of a one-bit quantizer, the power of the input to the correlator is constant, so there is no need to consider this effect.

While the first example given by Massey is the one depicted above, it was also shown that if the surrounding data is instead distributed as zero-mean Gaussian random variables, the optimum rule has precisely the same form as (4.1). Further, Massey demonstrated the optimum rule in an AWGN channel with essentially BPSK signaling and interestingly, subsequent studies in more complex settings have arrived at similar conclusions. In [57], it was shown that the optimum rule for location a synchronization sequence with an M-ary signalling scheme over the AWGN channel consisted of essentially the same components as (4.1) from [54]. They were also able to again demonstrate that the high SNR approximation performs very close to the optimum rule. Going further, AWGN channels with ISI was considered in [58], flat fading channels in [59] and frequency-selective channels in [60], to name some of the papers in this field. Similar conclusions was also reached by [53] and [46] under different sets of assumptions, where the maximum likelihood (ML) rules again and again was showed to involve a correlation in one way or another. All of these results have decisively demonstrated that the core principle of that finding the maximum in a correlation spectrum is central to timing synchronization.

Design of synchronization sequences

Regarding the design of synchronization sequence, it is intuitively clear that there are at least some designs which are outright unusable. Given that the task is to locate a known sequence somewhere within a received signal, we can generally say that the sequence must be unique. This means that if our synchronization sequence where to appear by chance somewhere else in the window, we will not be able to uniquely determine its position. However, if the sequence does appear twice in the received signal and the distance between the repetitions is known, we can use that information to temporally locate the data. We will soon return to this idea and for now focus on finding a single, known sequence within a window of received samples.

To avoid an instance where our synchronization sequence appears more than once by chance rather than by design, the sequence must be of sufficient length so that the probability of it showing up as a result of a stochastic process is sufficiently small. Further, as we established that the correlator is the optimum metric to maximize, we would like the sequence to have a peak at zero lag and, preferably, zero at all other timing instants. Due to the implementation specifics, we can place additional constraints of the sequence, such as the PAPR. Fundamentally, however, its correlation properties are paramount. Other important aspects of a synchronization sequence could be, as listed in [61], minimal overhead as well as rapid and low-complexity detection.

Early work in this field came from R.H. Barker, who developed the well-known Barker sequences [61]. These were originally formulated with a specific condition on the autocorrelation properties and the original solution was later slightly relaxed into what is now referred to a Barker sequences [62]. A sequence $\mathbf{s}_B[n] \in \{-1, 1\}$ of length L is a Barker sequence if its circular autocorrelation R^c is

$$|R^c_{\mathbf{s}_B}[m]| = \{0, 1\} \quad , \quad \forall\, m \neq 0 \tag{4.4}$$

where $R^c_{\mathbf{s}}$ is defined as [63]

$$R^c_{\mathbf{s}}(s_i, s_j) = \sum_{i=0}^{N-j-1} s_i\ s^*_{i+j} + \sum_{i=N-j}^{N-1} s_i\ s^*_{i+j-N} \quad , \ j = 0, 1, \ldots N \tag{4.5}$$

If $i = j$, then (4.5) reduces to $\sum_{i=0}^{N-1}\ s_i s_i^*$.

Only eight of these sequences have been found, the longest of length $L = 13$. It has not been formally proven that no longer sequences can be found, but it exists an overwhelming number of indicators supporting that claim and via computer simulations, no sequence of length $13 \leq L < 10^{22}$ have been found [64]. From a practical point of view, it means that no usable sequence longer than 13 exists, as 10^{22} bits reserved for synchronization is would certainly be infeasible. Barker sequences have had a number of uses over the years, for example to increasing the length of a radar signal via pulse compression [65].

Another type of sequence with specific autocorrelation properties is the Maximal Length Sequence, or m-sequence for short. It was originally developed by Solomon Golomb and has been extensively used since its inception in the 1950s [66]. They can be defined with the aid of linear feedback shift registers, where certain configurations will yield periodic sequences with the maximum possible period. In Figure 4.3, one such configuration is shown. The state vector $\mathbf{r}$ of this particular setting will, given that the initial state is not the zero vector, cycle through every possible five-bit combination except for the all-zero state exactly once before repeating itself. This means that the output vector $\mathbf{s}$ will be periodic with $2^5 - 1 = 31$ and in general, an linear shift register of degree N will produce a m-sequence of period $2^N - 1$ [67]. Note further that different initial values of the shift registers, apart from the zero vector, all produce cyclically shifted versions of the same sequence.

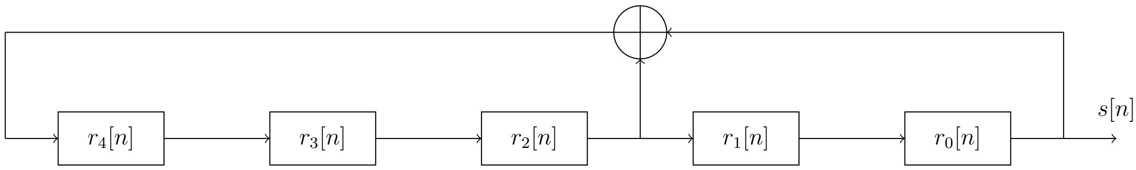

Figure 4.3: The generator for the m-sequence used in LTE [68], where the initial state is set to $\mathbf{r} = [0\ 0\ 0\ 0\ 1]$.

M-sequences have a number of interesting properties and the key characteristic yielding them especially useful for synchronization, relates to their autocorrelation properties. Feeding the binary vector $\mathbf{s}$ to a binary phase shift keying (BPSK) modulator, we produce the vector $\tilde{\mathbf{s}}$, i.e. $\tilde{s}[n] = 1 - 2s[n], \quad n \in \{0, 1, \ldots, N\}$. Then, the autocorrelation $C_{\tilde{s}}$ is given as [69]

$$C_{\tilde{s}}[m] = \sum_{n=0}^{N-1} \tilde{s}[n]\ \tilde{s}[(n+m)_N] = -1 \quad , \quad \forall\ m \neq 0 \tag{4.6}$$

Note that (4.6) is the same circular autocorrelation defined in (4.5), however expressed more compactly with the circular shift notation $(\cdot)_N$.

A key difference between the m-sequence and the Barker sequence, is that there is not any upper bound on the length of the m-sequence. This means that the difference between the main peak of the correlation spectrum and the side peaks can be made arbitrarily large. In Figure 4.4, this difference is shown. The peak of both spectra is equal to the length their respective sequence, so as the m-sequence can be designed to have any length, we can make the difference between the peaks significantly larger than for the longest possible Barker sequence.

A key aspect of the m-sequence is that within one period, it appears almost completely as white noise. The definition of white noise states that its autocorrelation at any locations other than the zero lag should be identically zero, and from Figure 4.4b, it is clear that the m-sequence approximates this to a stunning degree. The fact that an m-sequences is a completely deterministic sequence that appears fully random, is perhaps the major reason that they have found such widespread use. This is also why they are sometimes called pseudonoise (PN) sequences. They are for example used as a modulating sequence in code-division multiple access (CDMA) systems, as well as for synchronization purposes in LTE [70, 68].

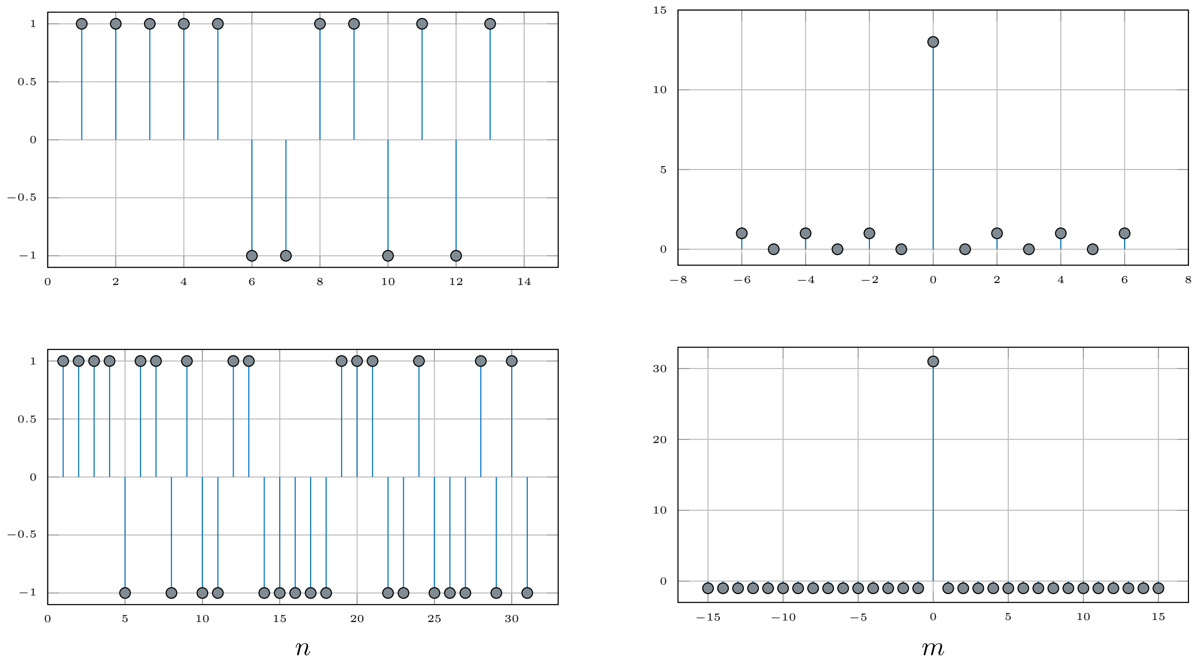

(a) The sequences. **(b)** The correlation spectra.

Figure 4.4: A length 13 Barker sequence (top) and a length 31 m-sequence (bottom) and their respective correlation spectra.

Lastly, we mention the topic of the cross-correlation of two different m-sequences. This is an important aspect in CDMA, where each user requires a unique sequence to modulate its data with and we would, preferably, like the cross-correlation of two different m-sequences to be close to zero, meaning that a m-sequence from one user appears as white noise when correlated with any other m-sequence used in the system. This leads us to the topic of Gold sequences, which we will not look into further, but note that it is possible to construct a set of sequences such that their respective cross-correlation is bounded. The construction of this set is based on the m-sequences and Gold sequences have found significant use in various contexts [66, 69].

We have now covered some the relevant background with regards to binary sequences. The concepts developed so far can be extended to cover sequences consisting of a larger alphabet, such as those made from a string of complex numbers. We will return to this topic when we discuss the synchronization procedure in LTE, as it is not applicable for any of the other methods that we will investigate.

Data-aided and non-data aided synchronization

The discussion so far has fundamentally been about our ability to detect a known sequence and the design of such a sequence. Going further, we could also question whether the synchronization sequence actually needs to be explicitly known at the receiver. This was hinted to earlier and for example, say that each payload was preceded by a repetition of some unknown sequence but of known length. In this setting, we could simply look for a place in the window where a given number samples of are repeated once or more times to accurately determine the start of the payload. With this approach, the received signal only needs to contain some predetermined structure, rather than a specific sequence. Broadly speaking, we can categorize synchronization methods into two major groups called data-aided (DA), where the sequence is explicitly known or non-data aided (NDA) methods, where some structural property is used. We will discuss methods of both types in following sections.

4.2 Cyclic prefix-based synchronization

As implied by title of this section, these methods relies on the cyclic prefix to find symbol timing and frequency offsets. Clearly, this is only relevant to systems using a CP, which OFDM happens to do. As mentioned in Section 2.1, we normally prepend each OFDM symbol with a prefix of length P made up of the last P samples the symbol. Consequently, an OFDM symbol itself already has repetitive structure, which can be used for synchronization. As the exact content of the cyclic prefix is unknown to the receiver, we would call this a NDA method.

An early paper in this field came in 1995 [71, 52], treating mainly the timing offset estimation. The same authors developed the concept further to cover both the frequency and timing offset estimates in a well-known paper from 1997 [53]. Additional developments can also be found in [72], where the idea is extend to synchronization in multiuser OFDM and in [73], where pilots used in channel estimation are exploited in tandem with the cyclic prefix to increase the accuracy of the offset estimation.

We will now present an outline of the method and its performance. Assume we transmit OFDM symbols with N subcarriers, each prepended with a cyclic prefix of length P over an AWGN channel. At the receiver, a window of length $2N+P$ is observed, as depicted in Figure 4.5. The received signal is affected by some unknown normalized frequency offset ϵ_0 and the correct location of the window is defined to be in the start of the cyclic prefix and parameter measuring the shift from this location is called μ. This means that the STO δ_0 as defined in Chapter 3 relates to μ as $\delta_0 = \mu - P$.

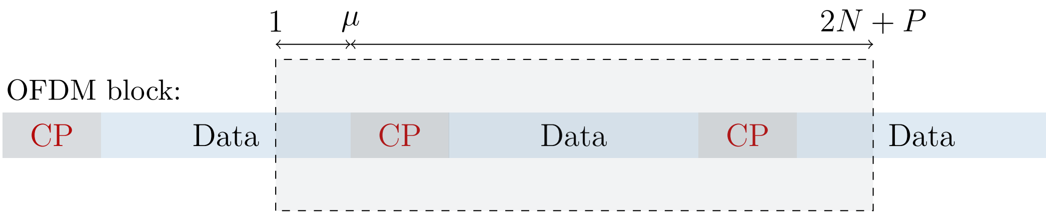

Figure 4.5: The window of received samples considered in [53].

We collect the received samples $y[n]$ in the vector $\mathbf{y} = \{y[1], y[2], \ldots, y[2N+P]\}$. The log-likelihood function $\mathcal{L}(\mu, \epsilon)$ of $p(\mathbf{y}|\mu, \epsilon)$ can then be shown to be [52, 53]

$$\mathcal{L}(\mu, \epsilon) = |\gamma(\mu)| \cos\left(2\pi\epsilon_0 + \angle\gamma(\mu)\right) - \rho\Phi(\mu) \tag{4.7}$$

where

$$\gamma(m) = \sum_{n=m}^{m+P-1} y[n]y^*[n+N] \tag{4.8}$$

$$\Phi(m) = \frac{1}{2}\sum_{n=m}^{m+P-1} |y[n]|^2 + |y[n+N]|^2 \tag{4.9}$$

$$\rho = \left| \frac{\mathbb{E}\Big\{y[n]y^*[n+N]\Big\}}{\sqrt{\mathbb{E}\Big\{|y[n]|^2\Big\}\mathbb{E}\Big\{|y[n+N]|^2\Big\}}} \right| = \frac{\text{SNR}}{\text{SNR}+1} \tag{4.10}$$

Maximizing (4.7) w.r.t. to ϵ_0 yields the following maximum likelihood (ML) estimate for ϵ

$$\epsilon_{\text{ML}}^{\text{est}} = -\frac{1}{2\pi}\angle\gamma(\mu) \tag{4.11}$$

assuming that $|\epsilon_0| < 1/2$. This is valid as long it can be assumed that a coarse synchronization in the acquisition step has already confined the CFO within this interval. With other methods of synchronization, this restriction can be lifted. In fact, as we will see, the range of the frequency estimator can be arbitrarily long with a suitable synchronization symbol.

We also write down the ML estimate for μ as

$$\mu_{\text{ML}}^{\text{est}} = \arg\max_{m}\ |\gamma(m)| - \rho\Phi(m) \tag{4.12}$$

In order to estimate both parameters, the estimate μ^{est} is then used as the argument to $\gamma(\cdot)$ in (4.11). Note that maximum likelihood estimator for the timing offset in (4.12) has the same components as (4.1), where the first term (4.8) is a correlation term and (4.9) a power correction term. In Figure 4.6, the log-likelihood function with respect to μ is shown. We look for the m that maximizes this function to find μ^{est}. In this figure, we have transmitted several OFDM symbols and we can clearly discern the peaks from each symbol.

The height of each peak compared to the other points of the correlation spectrum is tied to the length of the cyclic prefix. As we correlate with a longer known sequence, the likelihood of any random sequence of samples appearing similar decreases. This leads us to the question if performance gains can be achieved by simply extending the cyclic prefix beyond the length it is normally designed to be. This question was answered in [71], where they empirically demonstrated that there are no performance gains for STO estimation beyond a certain length, but there are for the CFO estimation. The exact length where the mean square error (MSE) behavior of the STO estimation ceases to improve is dependent on the SNR, which can be understood as locating a single peak in a noisy environment can never be more accurate that approximately the variance of the noise itself. For the CFO estimation, however, a longer cyclic prefix can be viewed as more points to average over, as the sum in (4.8) is made longer. Hence, the MSE behavior is improved for the frequency estimator as the length of the cyclic prefix is increased.

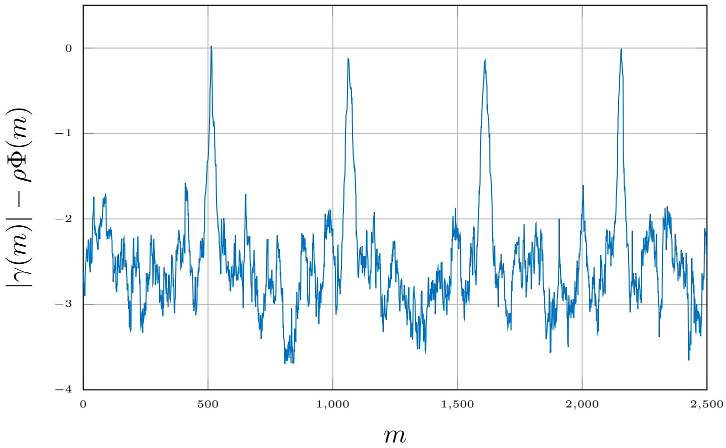

Figure 4.6: The log-likelihood function with respect to μ. Transmitted over an AWGN channel with 512 subcarriers, cyclic prefix of length 36 and SNR at 10 dB.

In [71], one-bit quantization was actually considered. The motivation was to design a synchronization method with very limited complexity, hence yielding affordable implementation. The paper only considers timing offset estimation and they showed that the maximum likelihood estimator $\mu_\mathrm{q}^\mathrm{est}$ of the true offset μ in the one-bit case can be written as

$$\mu_\mathrm{q}^\mathrm{est} = \arg\max_m \sum_{n=m}^{m+P-1} \Re\left\{r[n]r^*[n+N]\right\} \tag{4.13}$$

where $r[n] = \mathcal{Q}(y[n])$. Interestingly, this is highly reminiscent of S_1 in (4.1). This agrees well with our previous conclusion from Figure 4.2b, where we noted that with one-bit quantization, the correction term can be disregarded as the power of the quantized signal is constant.

The performance of this estimator is shown in Figure 4.7. Both Figure 4.7a and Figure 4.7b show the MSE of estimators. This is a common metric to measure the performance of an estimator and one that we will continue to use. In Figure 4.7a, we note that the MSE does not seem to decrease to below 10 before the SNR nears 4 dB in the infinite-precision case and around 10 dB in the one-bit case. This can be considered a lower limit on the usefulness of the method – if the MSE is larger than the cyclic prefix itself, it is clear that the estimates are highly uncertain. In Figure 4.7a, we also note that the saturation at low SNR a simulated artefact due to the finite length of the received OFDM blocks.

In Figure 4.7b, we have in addition to the MSE for the frequency estimator also plotted the Cramer-Rao bound [51]. It is given as

$$\mathrm{CRB} = \frac{1}{2\pi^2}\,\frac{3\ \mathrm{SNR}^{-1}}{N(1-1/N^2)} \tag{4.14}$$

With any estimation process, the accuracy at a given SNR is bounded by a fundamental limit.

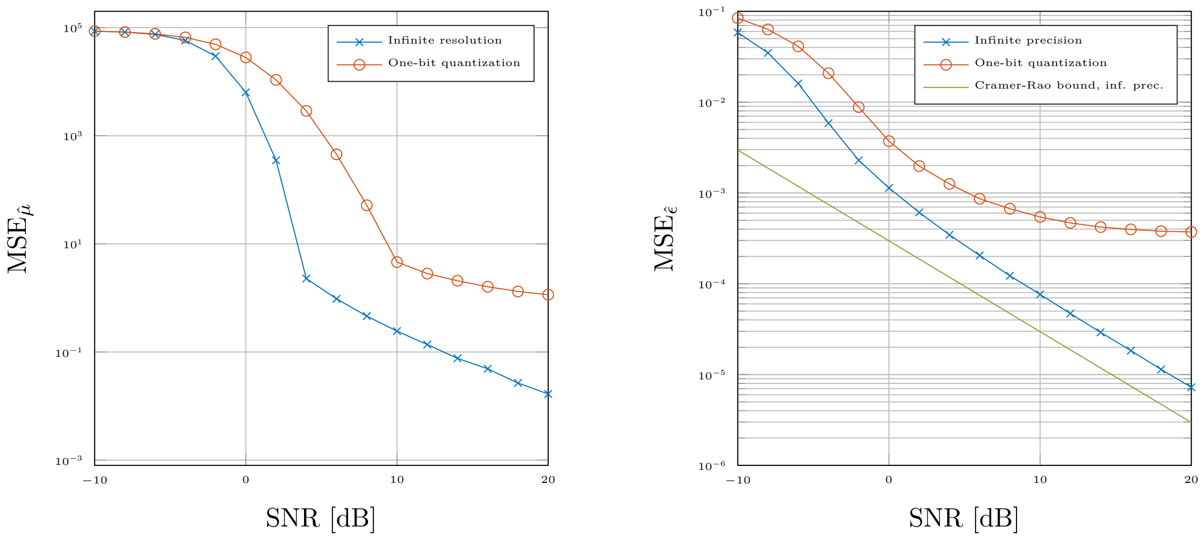

(a) The MSE of the timing estimator.

(b) The MSE of the frequency estimator, $\epsilon_0 = 0.2$.

Figure 4.7: The MSE of the timing and frequency estimators in both the infinite-precision and one-bit quantization case. Results obtained from simulation over an AWGN channel, 512 subcarriers, cyclic prefix of length 36 and $5 \cdot 10^4$ trials at each SNR.

These bounds are a consequence of a concept central in statistical inference known as the Fisher information. On an intuitive level, this quantity captures the shape of the likelihood function and tell us how much we can infer about a parameter from an observation depending on that parameter. It is a useful tool when evaluating different estimators, as it provides a reference to the optimal behavior of an estimator.

As we see in Figure 4.7b, the MSE in the one-bit case seems to have a floor. This is due to the quantization distortion, which does not decrease with reduced thermal noise power. We can finally note that in the SNR regime below 0 dB, the performance in the one-bit case seems to be limited by thermal noise and above that, the quantization distortion starts to dominate.

An important aspect of synchronization algorithms is to consider their robustness, meaning their ability to distinguish the synchronization sequence from noise. The two interesting quantities is the missed synchronization probability, meaning how likely the method is to miss that a synchronization sequence is present and the false detecting probability, meaning how likely the method is to mistake noise for an actual signal. By recording the values of the likelihood functions

$$\mathcal{L}_{\mathrm{inf}}(m) = |\gamma(m)| - \rho\Phi(m) \tag{4.15}$$

in the infinite-precision case and

$$\mathcal{L}_{\mathrm{q}}(m) = \sum_{n=m}^{m+P-1} \Re\left\{r[n]r^*[n+N]\right\} \tag{4.16}$$

in the quantized case at $m = \mu$, both when there is a synchronization sequence present and when there is not, we can produce Figure 4.8.

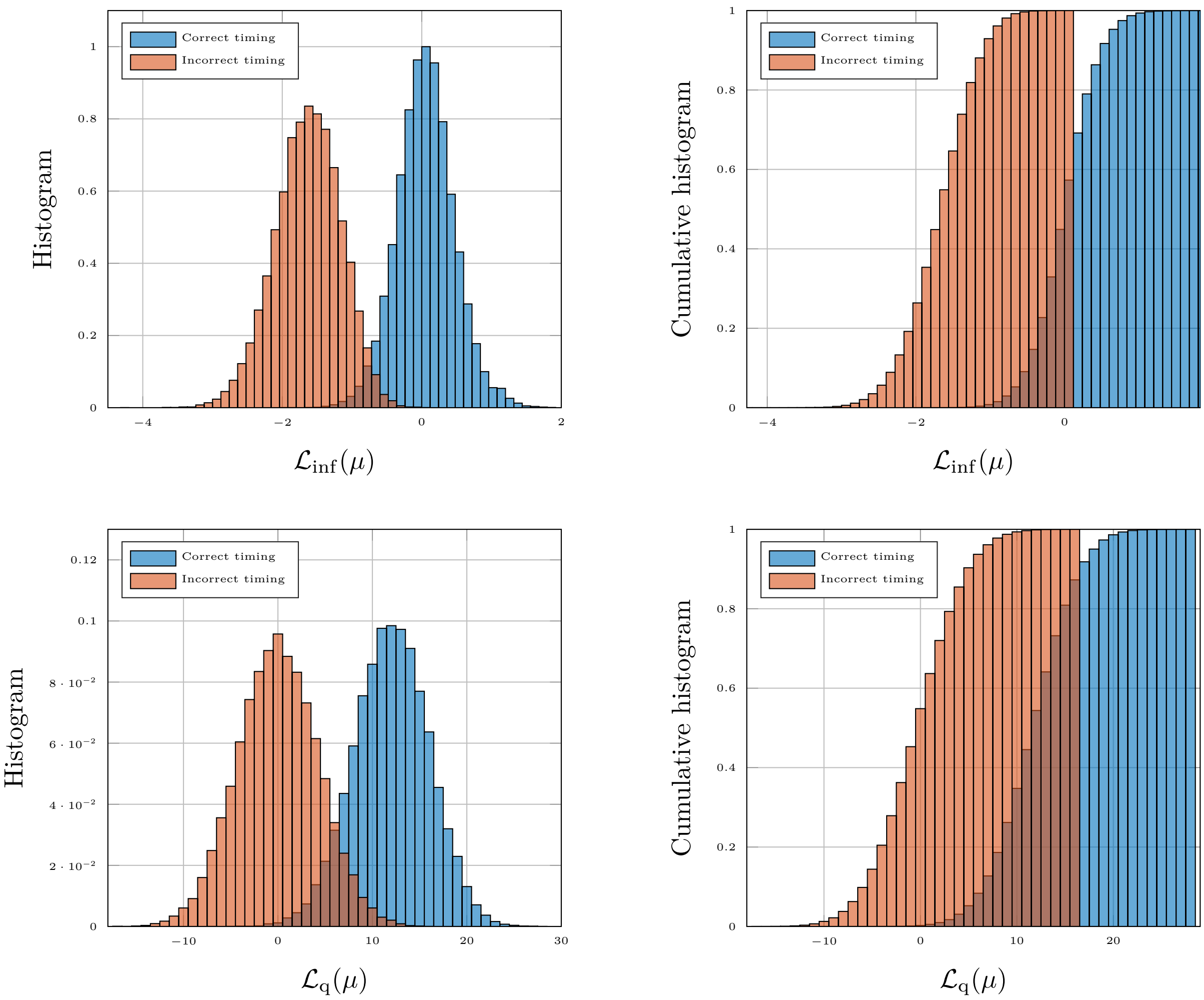

(a) Normalized histograms. **(b)** Normalized cumulative histograms.

Figure 4.8: Histogram and cumulative histogram of infinite-precision estimator (top) and one-bit estimator (bottom). Results obtained via recording $5 \cdot 10^4$ values of each estimator when there was a signal and when there was not. SNR set to 0 dB.

In Figure 4.8a, we see the histograms of the value of the estimator both when there was a signal present and when there is not. We see that clearly, the average value of (4.15) and (4.16) is larger when the signal is present than when it is not, but the overlap indicates that the probability of mistaking the two cases is not negligible. As expected, the overlap is more significant in the one-bit case, which agrees with all previous results. We can also sum the histogram in a different way to produce Figure 4.8b, where the cumulative histograms are shown. We can view these as an approximation of the true cumulative distribution of (4.15) and (4.16). If we denote these approximations with $\tilde{P}^{\text{c.}}$ and $\tilde{P}^{\text{inc.}}$ for the correct and incorrect timing in the infinite-precision case respectively, we find the missed synchronization probability (MSP) and the false detection probability (FDP) at some threshold τ as

$$\mathrm{MSP}_{\mathrm{inf}}[\tau] = \tilde{P}^{\mathrm{c.}}\left[\mathcal{L}_{\mathrm{inf}} < \tau\right] \tag{4.17}$$
$$\mathrm{FDP}_{\mathrm{inf}}[\tau] = 1 - \tilde{P}^{\mathrm{inc.}}\left[\mathcal{L}_{\mathrm{inf}} < \tau\right] \tag{4.18}$$

Constructing the MSP and the FDP for the quantized case in the same manner, we can then match the thresholds and produce Figure 4.9.

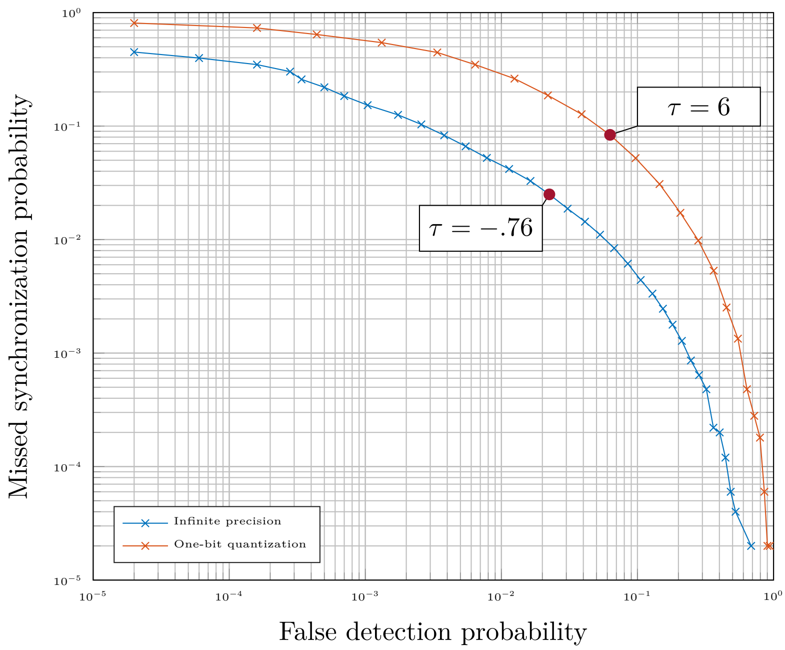

Figure 4.9: The relationship between the missed synchronization probability and the false detection probability when using the cyclic prefix-based synchronization method.

We see in Figure 4.9 that as one parameter is increasing, the other is decreasing. This is understood as we set the threshold higher and higher, the probability of mistaking noise for the actual signal tends towards zero, but at the same time, we increase the risk of missing the signal when it is present. From the figure, we note that it is possible to set threshold so that both the MSP and FDP is below 3% (infinite-precision) and 10% (one-bit).

As a last note before moving on, we note that it is mentioned in the original paper that the accuracy can be greatly improved by averaging over multiple cyclic prefixes. By using synchronized receive antennas and assuming that the signal arrives at the same time at each antenna, we can combine the likelihood functions from each antenna to produce a better estimate. The results are depicted in Figure 4.10.

We see that that in the average of 10 antennas, a clear peak is visible in both Figure 4.10a and Figure 4.10b. When examining the peaks of all the colored lines, we note that it would not have been possible to use this method of synchronization at this SNR level in a single-antenna system. As a last note on this topic, the multi-antenna gains in a more realistic channel model would probably be less pronounced. The reason that the simple averaging method works so well, is that

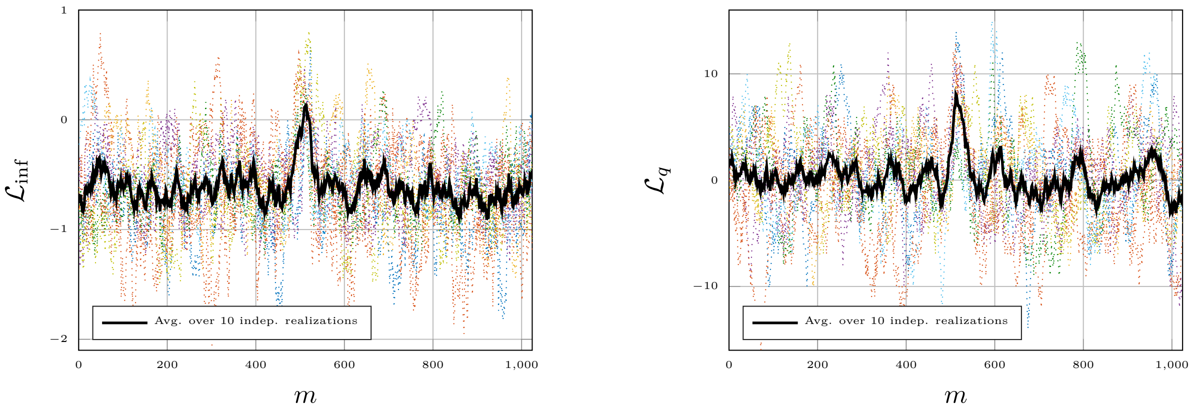

(a) Infinite-precision. **(b)** One-bit quantization.

Figure 4.10: The likelihood functions for the symbol timing estimators. The colored lines are the results on each antenna and the black line is their average. Results obtained with 512 subcarriers, cyclic prefix of length 36 and an all-ones channel with an SNR per antenna set to −4 dB.

the noise on each antenna is assumed to be fully independent. In a fading MU environment, this assumption is not valid and consequently, the correlated distortion would not be averaged out. Moreover, expanding the model to include for example hardware imperfections in the receiver chain also introduces distortion that is correlated across all receive antennas, which would further limit the performance of simple averaging.

As we have just seen in the previous section, using the periodicity of a cyclic prefix is a viable way to extract information of the STO and CFO. However, there are some shortcomings. For example, the results in Figure 4.7 are for the AWGN channel – if the channel model would instead have been a more realistic multipath fading model, the cyclic prefix would clearly have been affected by the channel delay spread. The consequence of that would be that the correlation decreases further, providing additional difficulties in locating the peak. Further, the gap between the frequency estimator and its Cramer-Rao bound is roughly 4 dB in the infinite-precision case, perhaps hinting towards possibilities for improvement. A viable path forward would be to consider a dedicated training symbol, instead of relying solely on the repetition present in the OFDM symbol itself. This leads us to our next method, originally presented by Schimdl and Cox in the mid-1990s.

4.3 Schmidl and Cox synchronization

Around the same time that the method described in the previous section was developed, other ideas were investigated by other groups. For example, it was suggested in [46] that a repeated OFDM symbol could be used. The same idea was present in works of other groups as well and building upon that earlier work, Schmidl and Cox [48] produced a paper that have since received widespread attention. They were able to expand and reduce the complexity of the earlier work, resulting in a simple method with surprisingly good performance.

The full version of the method in [48] uses two dedicated OFDM symbols. The first is mirrored across the middle, meaning that the first half of the symbol is identical to the second half. We can construct a symbol like this by putting data on every other subcarrier and leaving the rest as null. In general, we can construct a symbol with q repetitions by modulating every qth subcarrier and leaving the rest as zeros. In the Schmidl and Cox method, a PN sequence is used to generate QPSK symbols which are placed on the even subcarriers, while zeros are placed on the odd. This is depicted in Figure 4.11a. As symbols are only transmitted on half of the available subcarriers, these symbols are scaled with $\sqrt{2}$ to make the energy constant over the full OFDM frame. We collect all these symbols in the vector $\mathbf{c}_1$.

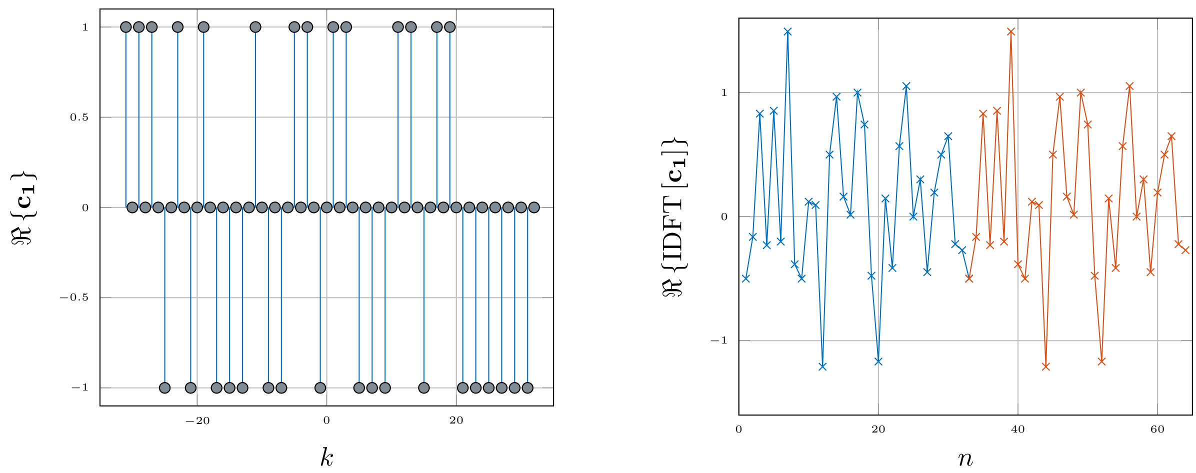

(a) The real part of the symbols in $\mathbf{c_1}$

(b) The real part of the first synchonization symbol.

Figure 4.11: The first synchronization symbol in the method by Schmidl and Cox.

In Figure 4.11b, we show the first OFDM symbol used in this method. We see that the symbol is indeed constructed from two equal parts, which we will use to get an estimate for the timing information as well as for the fractional CFO. As we saw in Chapter 3, the rotation of sample is a function of its index n. This means that if we have two identical symbols arriving at time indices $n_1, n_2,\ n_2 > n_1$, we know that their phase difference will be proportional to $\Delta n = n_2 - n_1$. If we are only comparing the phase of two points on the unit circle, we can only determine their difference $\Delta\varphi$ within the interval $\varphi < |\pi|$. Expressing their phase difference relative to the subcarrier spacing, repeating a sequence once within an OFDM symbol will limit the acquisition range to $|1/2|$ subcarrier spacing, i.e. the estimated CFO $\epsilon^{\text{est}} < |1/2|$.

To find larger frequency offset, the second OFDM symbol is used. This symbol consists of a new random sequence of QPSK symbols on its even subcarriers and another on its odd. We collect all of the symbols on the even subcarriers in the vector $\mathbf{c}_2$. Next, we record the phase difference between the symbols in $\mathbf{c}_1$ and $\mathbf{c}_2$ in $\nu[k] = \sqrt{2}c_2[k]/c_1[k],\ \nu[k] \in \{\pm 1, \pm j\}$. This vector is assumed to be known at the receiver. As mentioned in in Chapter 3, integer frequency offset will appear as a shift of the DFT output in the receiver. By determining the number of positions the output of the DFT is shifted from its expected position, we can determine the integer frequency offset. The sequence on the odd subcarriers are not used specifically in the original paper, but it can for example be used as pilot symbols to estimate the channel on the odd subcarriers.

To find an estimate of the STO, Schmidl and Cox suggests the following metric

$$M[m] = \frac{|P(m)|^2}{(R[m])^2} \tag{4.19}$$

where

$$P[m] = \sum_{n=0}^{N/2-1} y[m+n]y^*[m+N/2+n] \tag{4.20}$$

$$R[m] = \sum_{n=0}^{N/2-1} |y[m+N/2+n]|^2 \tag{4.21}$$

We see that this is simply a normalized correlation and the timing offset in found by detecting the peak in M, $\delta^{\text{est}} = \arg\max_m M[m]$. Note that we now use the letter δ_0 to again denote the timing offset, as this method will find on the start of that *data* in the symbol instead of the start of the symbol itself.

The timing metric M is shown in Figure 4.12a.

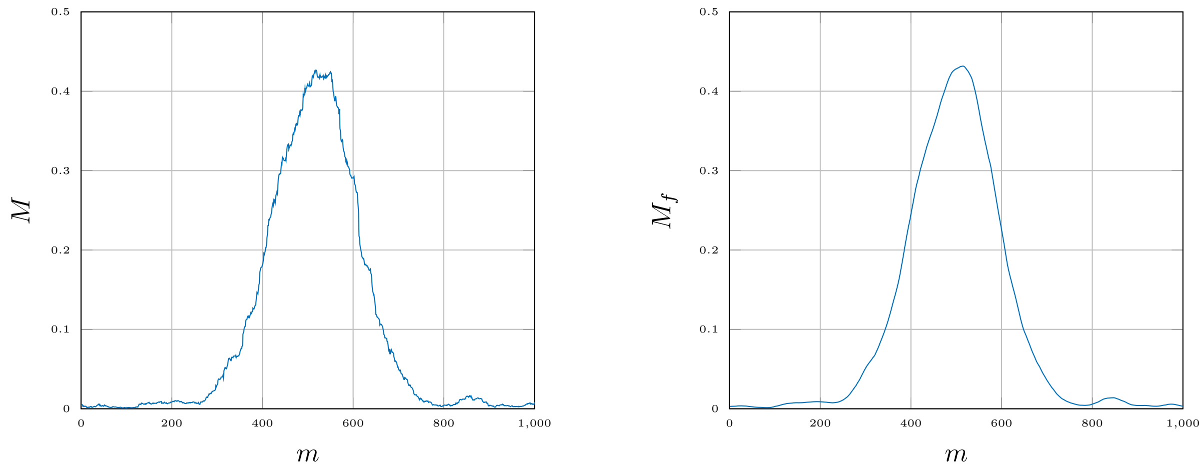

(a) The timing metric M.

(b) The timing metric M_f.

Figure 4.12: The timing metric proposed by Schmidl and Cox as well as the metric proposed by Minn [74]. Obtained with 512 subcarriers, cyclic prefix of length 36 in a frequency-flat channel model with SNR set to 10 dB.

Due to the presence of the cyclic prefix, the timing metric exhibits a plateau. The width of this part will be as wide as the cyclic prefix P in the AWGN and in the flat-fading channel, and roughly $P - L$ wide in a frequency-fading channel of length L. As long as the starting point is taken at any point within the plateau, no ISI or ICI is incurred, as we showed in Chapter 3. The rotation caused by showing the start of the symbol to be somewhere in the cyclic prefix can be compensated in the channel equalizer using a single pilot. However, to reduce the uncertainty slightly, the following metric was proposed in [74]

$$M_f[m] = \frac{|P[m]|^2}{[R_f[m]]^2} \tag{4.22}$$

where R_f is given by

$$R_f[m] = \frac{1}{2} \sum_{n=0}^{N-1} |y[m+n]|^2 \tag{4.23}$$

We see that R_f is normalized with the half the power of the entire symbol instead of all of the power of half the symbol, as in (4.21). The metric M_f is then averaged over a window of size $P+1$, i.e.

$$M_{\text{avg}}[m] = \frac{1}{P+1} \sum_{n=0}^{P} M_f[m+n] \tag{4.24}$$

The estimate for the timing offset δ^{est} is then found as the maximizing argument to $M_{\text{avg}}[m]$.

The metric M_f is shown in Figure 4.12b. We see that M_f is essentially a low-pass filtered version of M, meaning that any high-frequency peaks present in M that may cause an error is smoothed out.

In Figure 4.13, the mean value and variance of M and M_f is shown. In Figure 4.13a, we see that M indeed follows the theoretical values predicted by the analysis in [48]. Further, we see that the mean of M_f is similar to that of M. This agrees with what we expected, as M_f is normalized in the same way as M. In the one-bit case, the results are similar, but due to quantization distortion, the two parts of the synchronization symbol are less correlated and, consequently, the mean values of the timing metric is lower. The effect is identical regardless of whether we use M or M_f which, for the same reasons as in the infinite-precision case, is to be expected.

In Figure 4.13b, the variance of the timing metric is shown. We again see that the analysis from [48] matches well with the simulations of M and that the one-bit case again displays reduced performance. Here, the purpose of M_f should become clear, as this metric has significantly lower variance than M. Via the averaging in (4.22), the metric is less sensitive to fluctuations, which explains why the variance is lower. Lastly, we note that the variance of the timing metric below $2 \cdot 10^{-3}$ for the entire range of SNR in both cases. A visual comparison to Figure 4.8a, where the variance at 0 dB is above 10^{-1} and 10^{1} in the infinite-precision and the one-bit case, respectively. This is indicative of the robustness of this method, which we will soon look at more in-depth.

In Figure 4.14, the MSE of the timing and estimators are shown. In Figure 4.14a, we note that the performance is similar or perhaps slightly worse than for the cyclic prefix method. Clearly, the strength of this method does not lie in MSE performance of the timing estimator. A reason for this is the plateau of the timing metric – even a low noise power may shift the maximum a

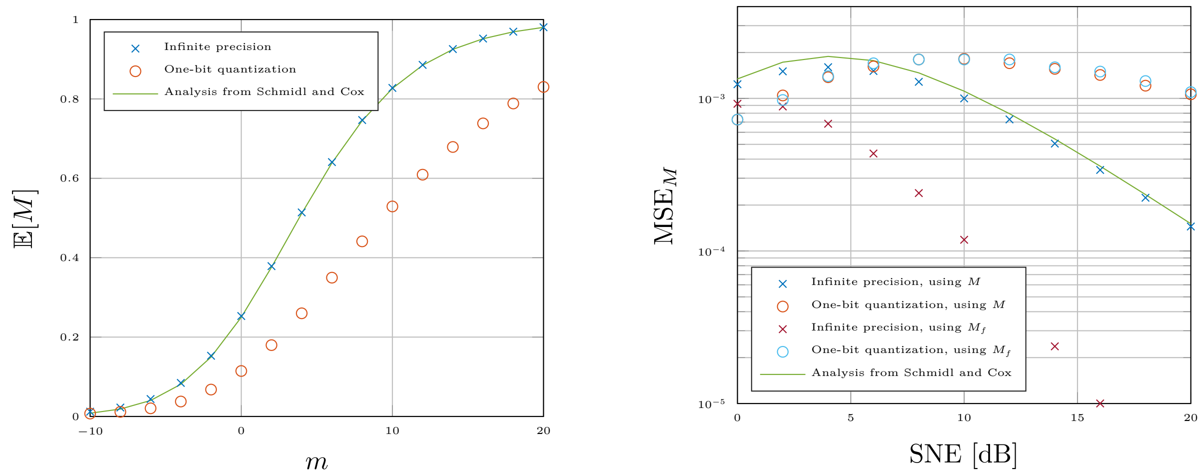

(a) The mean value of the timing metric M.

(b) The variance of the timing metric M.

Figure 4.13: The mean and variance of the timing metric M. Obtained via $5 \cdot 10^3$ simulations over an AWGN channel with 512 subcarriers and a cyclic prefix of length 36.

single step in either direction, which adds to the MSE. In [74], other methods where the training symbols are redesigned to yield sharper peaks is shown to outperform the results in Figure 4.14a. However, as any estimate δ^{est} that lies within the cyclic prefix does not add ISI, SNR levels around 5 dB in the infinite-precision case and roughly 10 dB in the one-bit quantization case is sufficient to obtain a decent estimate. However, in a mmWave scenario, the received SNR per antenna is likely to be well below 10 dB. The results therefore indicates that timing estimation with this method in a system with one-bit quantization might be difficult without some joint processing, using the signal from all or some of the receive antennas. Lastly, the saturation at low SNR again stems from simulation artefacts rather than actual results.

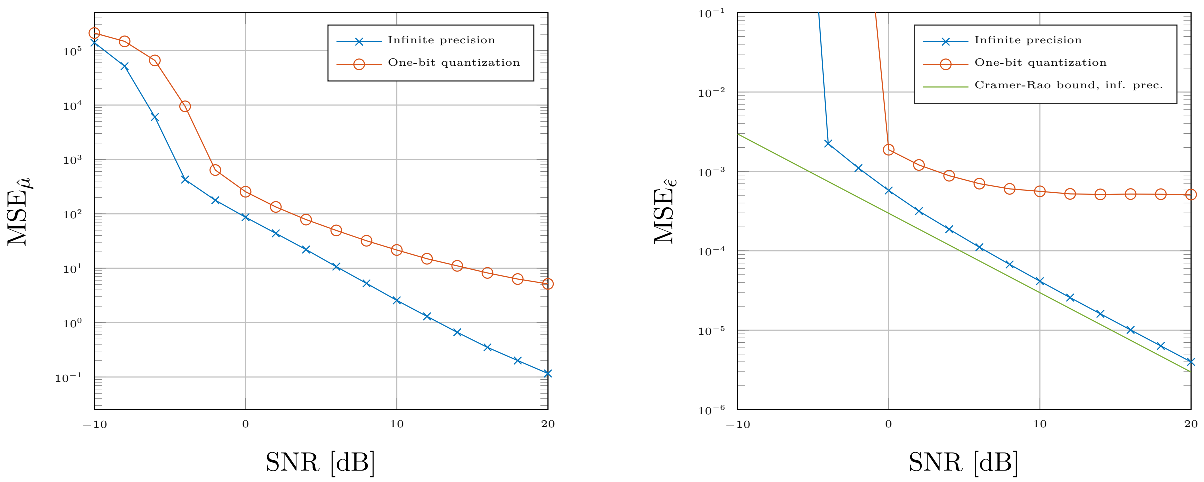

(a) The MSE of the timing estimator.

(b) The MSE of the frequency estimator, $\epsilon_0 = 2.2$.

Figure 4.14: The MSE of the timing and frequency offset estimators, using M_f as the timing metric. Obtained via $5 \cdot 10^3$ simulations over an AWGN channel with 512 subcarriers and a cyclic prefix of length 36.

In Figure 4.14b, we can see a clear improvement over the results obtained via the cyclic prefix method in 4.7b, at least in the infinite-precision case – the one-bit case has roughly the same performance as the cyclic-prefix method. The infinite-precision case has gained roughly 3 dB and is now only about 1 dB away from the Cramer-Rao bound. In [48], an analytic expression for the variance of ϵ^{est} is given as

$$\text{var}[\epsilon^{\text{est}}] = \frac{2}{\pi^2 \cdot N \cdot \text{SNR}} \tag{4.25}$$

We see that Figure 4.14b agrees well with the analysis.

At SNR below -4 dB in the infinite-precision case and below 0 dB, we see that the frequency estimator seems to completely break down. This is perhaps surprising, given that this effect is not visible in 4.7b and that for higher values of SNR, the Schmidl and Cox estimator outperforms the former. This is however due to the fact that the range of the cyclic prefix method is limited to $|\epsilon_0| < 1/2$, whereas the Schmidl and Cox method can estimate an arbitrarily large integer frequency offset as well. Hence, the potential error magnitude is virtually unlimited in this method compared to the cyclic prefix-based method, yielding vastly different MSE behavior. This again points to the fact that these methods can be difficult to use in a low SNR regime without joint processing in the receiver.

By collecting the values of M at the optimal position when there is a signal present and when there is not, we can produce Figure 4.15, similar to Figure 4.8.

We see in both Figure 4.15a and 4.15b that the is almost no overlap between the value of the timing metric when there is signal present and when there is not. This means that it is straightforward to set a threshold from frame detection without risking an abundance of false positives. Note that this does not mean that the Schmidl and Cox method produces more accurate timing estimates than the cyclic prefix method, only that it is better at differentiating between when there is a signal present and when there is not. Plotting the histograms together, we obtain Figure 4.16.

In 4.16a, we see that probability of observing a value of M larger than 0.1 is vanishingly small. We also see an approximation of the M when there is no signal present from [48] and that the simulation agree fairly well. In Figure 4.16b, we can see some overlap around 0.05, but otherwise it is largely the same trend. Figure 4.16 is highly illustrative of the main strength of the method develop my Schmidl and Cox – its robustness. The input to the correlator is two consecutive block of length $N/2$ and it is clearly highly unlikely that a similar structure to the training symbol would appear as a result of a random process.

Before moving on from the Schmidl and Cox method, we present some findings unique to the one-bit case. In Figure 4.17, the frequency dependence of the timing metric and, consequently, the timing and frequency estimators is illustrated.

We see in Figure 4.17a that the value of M at the optimal location seems to plateau for values of $\epsilon_0 \neq 0$. A similar effect is present in Figure 4.17c, where the case $\epsilon_0 = 0$ is yields significantly better MSE performance. For the MSE of the timing estimator in Figure 4.17b, it is not very clear

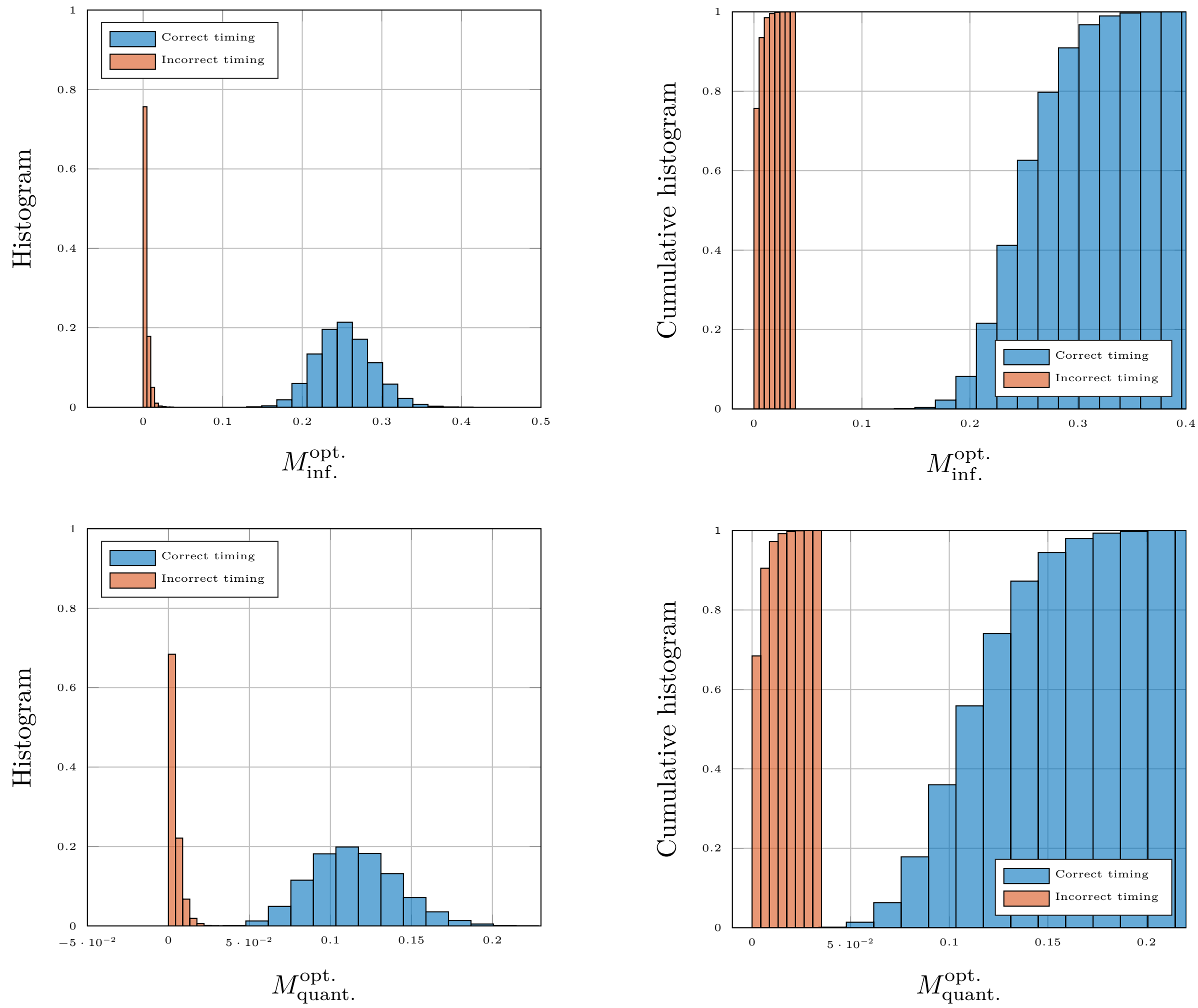

(a) Normalized histograms. **(b)** Normalized cumulative histograms.

Figure 4.15: Histogram and cumulative histogram of infinite-precision estimator (top) and one-bit estimator (bottom). Results obtained via recording $5 \cdot 10^4$ values of each estimator when there was a signal and when there was not. SNR set to 0 dB.

exactly how the CFO affects the performance, but clearly, the performance is not independent of the carrier frequency offset ϵ_0.

Intuitively, we can understand this as a frequency offset means that the received signal has been subject some rotation with respect to the transmitted signal. In the one-bit quantization case, this will cause some samples of the signal to move into another quadrant in the complex plane and, thus being quantized to another point. This can have significant effect, evident in Figure 4.17, as it will affect the similarity of the two halves of the symbols and thereby their correlation.

Lastly, we note that just as in the case of the cyclic prefix method, the performance can greatly benefit from joint processing in a situation with multiple receive antennas. With the similar assumptions as in the cyclic prefix case, we see the result in Figure 4.18.

The effect here is, as we would expect, an improvement of the MSE performance of the timing

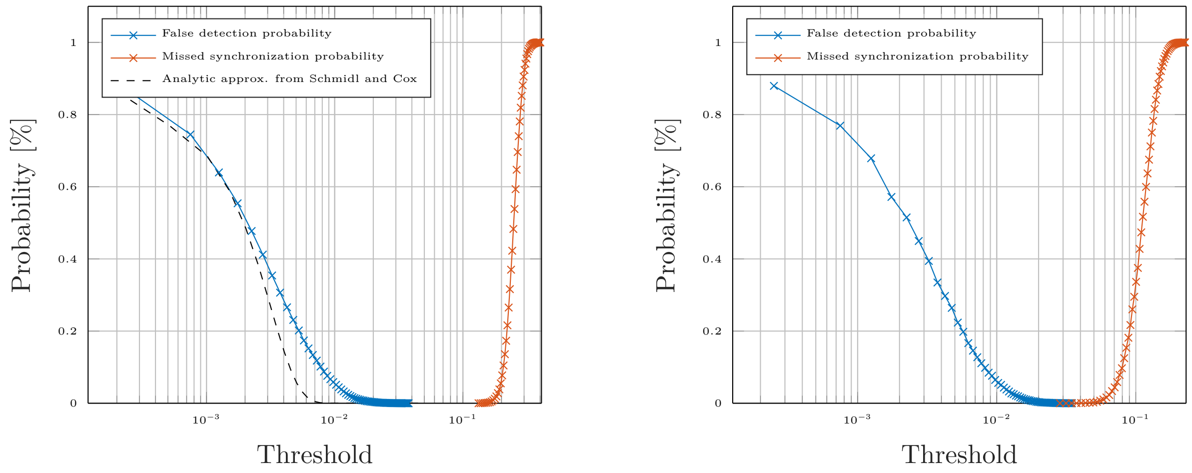

(a) Infinite-precision case. **(b)** One-bit quantization case.

Figure 4.16: The normalized histogram and cumulative histograms in the infinite-precision and one-bit quantization case. Results obtained via recording $5 \cdot 10^4$ values of each estimator when there was a signal and when there was not. SNR set to 0 dB.

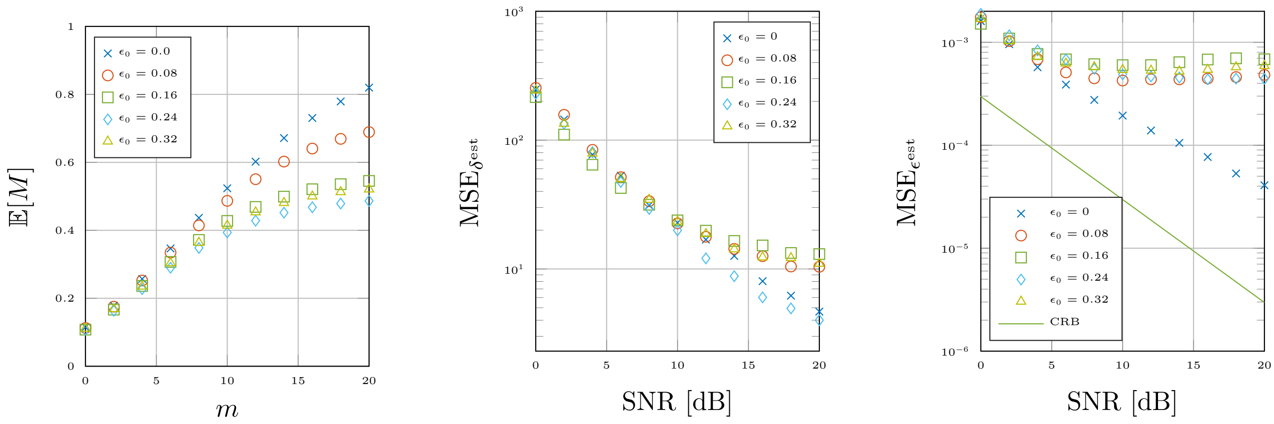

(a) Timing metric. **(b)** Timing estimator. **(c)** Frequency estimator.

Figure 4.17: Illustration of the frequency dependency of the estimator in the one-bit quantization case.

estimator. As mentioned previously, this type of performance gains requires that a number of assumptions are true, such as that the STO is the same on all antennas. Whether or not those are reasonable is another topic altogether, so we will here simply state that if it is, it would bring significant performance gains. From Figure 4.18, we note that 10 antennas on the receiving side is sufficient for the one-bit system to rival the performance of the infinite-precision system. This is an interesting finding, as it adds weight to the statement that moving from SISO systems to SIMO or MIMO will enable the use of low-resolution analog-to-digital converters. Even though this is a highly idealized case, it serves as motivation for further study into this topic.

The method presented by Schmidl and Cox provided several advantages over previously described methods. However, as seen in Figure 4.14b, the frequency estimation has a gap to the Cramer-Rao bound. This indicates that at least from a theoretical standpoint, a more accurate method could be devised. Moreover, and perhaps the most important drawback of the Schmidl and Cox method, is

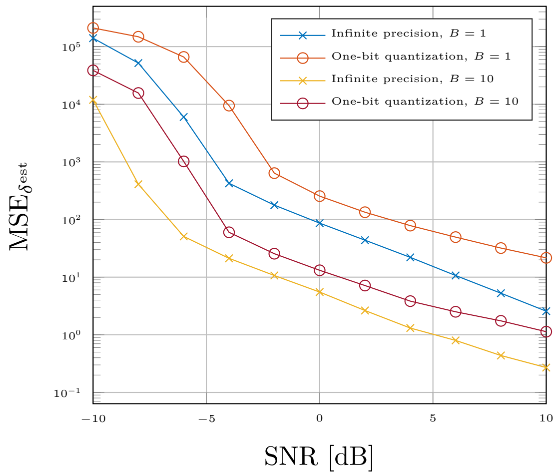

Figure 4.18: The MSE of the STO estimation with B antennas that can combine their correlation spectrums before coming to a decision. Obtained over the all-ones channel with 256 subcarriers.

that it requires two OFDM symbols in order to estimate and CFO larger than $\pm 1/2\Delta f$. Reducing this overhead was the major motivator in the development of an extension of this method, which will now be discussed in detail.

4.3.1 Morelli and Mengali extension

In 1999, a CFO estimation method was presented by Morelli and Mengali in [51]. In this paper, it was suggested to use a single training symbol consisting of q identical parts. In Figure 4.19, an example where $q = 8$ is shown.

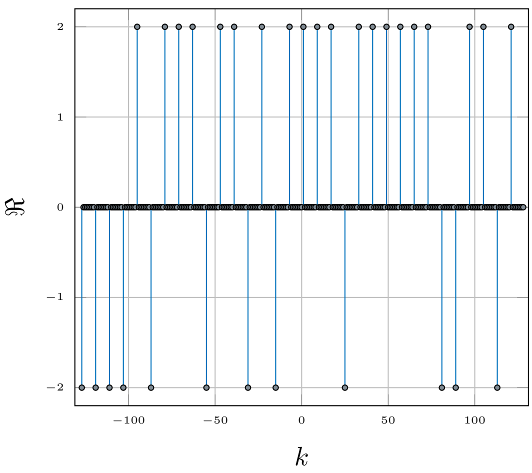

(a) The real part of the symbol

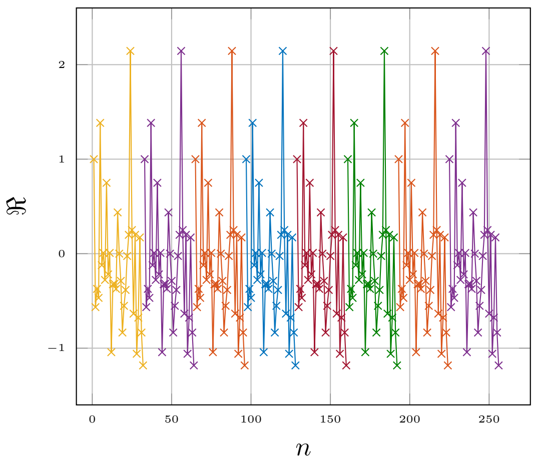

(b) The first synchonization symbol.

Figure 4.19: The real part of the first synchronization symbol in the method by Morelli and Mengali.

As mentioned in the previous section, we can construct these symbols by putting data on every qth subcarrier and leaving the rest as null, depicted in Figure 4.19a. In Figure 4.19b, we see the real part of the resulting symbol after the IDFT operation.

On the receiving side, Morelli and Mengali assumes that an accurate timing estimate has already been computed and, similarly to the Schmidl and Cox method, consider the angle of the correlation output. The key difference here is that in [51], they have q angles to consider instead of a single angle, as in the Schmidl and Cox method. The correlator output is given by

$$\varphi[m] = \sum_{n=1}^{\hat{\delta}+N-1-mN/q} y[n-mN/q]y^*[n], \quad m = 1, 2, \ldots, q/2 \tag{4.26}$$

Note that in (4.26), we use the estimate of the STO δ^{est}. This can be obtained in a similar manner as in the Schmidl and Cox method, given that q is a multiple of two.

Having obtained these angles, Morelli and Mengali construct the best linear unbiased estimator (BLUE) of the parameter ϵ_0 as

$$\epsilon^{\text{est}} = \frac{1}{2\pi/q} \sum_{m=1}^{q/2} \chi[m] \arg\{\varphi[m]\varphi^*[m-1]\} \tag{4.27}$$

In (4.27), we use χ to denote the vector of optimal weights. It is calculated in the original papers (cf. (16) in [51]), and given by

$$\chi[m] = 3\frac{(q-m)(q-m+1)-(q/2)^2}{\frac{q}{2}(q^2-1)} \tag{4.28}$$

For a BLUE estimator, its variance is known. With the parameters in this method, it is shown to be

$$\text{var}[\epsilon^{\text{est}}] = \frac{3(\text{SNR})^{-1}}{2\pi^2 \cdot N(1-1/q^2)} \tag{4.29}$$

We note two thing with the result in (4.29). Firstly, if we set $q = 2$, it becomes (4.25). This shows that the Morelli and Mengali method is indeed an extension of the original method by Schmidl and Cox, which was also pointed out in the original paper. Secondly, we note that as $q \to N$,(4.29) will approach the Cramer-Rao bound, given by (4.14). Consequently, as more repetitions are used, the variance of the frequency estimator will approach the lower bound. However, we cannot set q arbitrarily close to N, as with more repetitions, the length of each repetition decreases. With short repetitions, the likelihood of surrounding samples looking similar to our synchronization symbol increases, thereby increasing the risk of an erroneous estimate. For a given number of subcarriers N used in a system, the number of repetitions q should be set according to some specific system requirements.

In Figure 4.20, we show some simulation results with the Morelli and Mengali method.

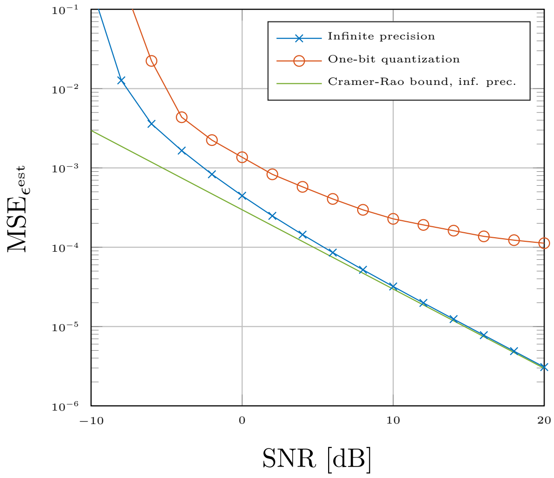

Figure 4.20: The MSE of the frequency estimator, $\epsilon_0 = 2.2$. Results obtained via $5 \cdot 10^3$ simulations over an AWGN channel with 512 subcarriers and $q = 8$.

In the infinite-precision case, the method shows improved MSE performance over the Schmidl and Cox method. As the SNR increases, the MSE seems to be approaching the Cramer-Rao bound, which, from (4.29) is no surprise.

In the one-bit case, we also see a slight improvement over the results from the Schmidl and Cox method. As the Morelli and Mengali algorithms essentially obtains several estimates of the CFO and combines them via (4.27), there is no obvious reason why the one-bit case should not benefit in a similar way. Still, we note that the MSE seems to level out as the SNR increases, since the quantized signal is vitiated with additional quantization distortion.

We have now examined the Schmidl and Cox method, as well as the extension provided my Morelli and Mengali. The latter of this two, while being slightly more computationally complex, is able to reduce the MSE of the CFO estimation as well as reducing the overhead by 50 %. Further, we noted that these methods outperform the cyclic prefix-based method in terms of robustness and accuracy of the CFO estimation. As both the Schmidl and Cox and Morelli and Mengali methods have used dedicated training symbols, they can be regarded as DA methods. These are more complex methods and require some aspects of the system to be predetermined. However, examining the results of our investigation, the performance gains are significant and the additional complexity might be small price to pay. Lastly, we will look at the performance of a method more akin to what is used in current 4G networks.

4.4 LTE-like synchronization

In current mobile network, commonly referred to by the name of the technology standard LTE, synchronization is achieved with two dedicated signals. These are called the primary synchronization signal (PSS) and the secondary synchronization signal (SSS). As the synchronization in LTE needs to accomplish more things than simply estimating the STO and CFO, we will gloss over the details and simply focus on the parts we are interested in. The interested reader is referred for example to [28], where more information can be found. Additionally, we note that the LTE standard does not specify exactly how to find the STO and CFO, only the signals that are available. This is the reason why we refer to this section as LTE-like, rather than anything else – the details of the methods in use differ from vendor to vendor and is generally not known to the public.

The PSS is used to obtain it initial timing estimate in LTE. For that reason, the problem of locating it is similar to the scenario we have considered in the previous methods and, hence, meaningful to compare. The PSS is a predetermined sequence that the receiver tries to locate via a correlator, so the basic steps are familiar. The exact construction of the PSS, however, is via a Zadoff-Chu sequence. These sequences, sometime referred to as Frank-Zadoff-Chu are an example of a constant-amplitude, zero autocorrelation (CAZAC) sequence and were developed in [75, 76, 63]. As mentioned in our initial discussion on the design of synchronization sequences, the ideas of designing a sequence with good correlation properties can be extended to non-binary sequences, of which the Zadoff-Chu sequence is an example. The sequence is set to length 63 and is constructed in the frequency-domain as [68]

$$s_{\mathrm{ZC}}[n] = \begin{cases} \mathrm{e}^{-j\frac{\pi u n(n+1)}{63}}, & 0 \leq n \leq 30 \\ \mathrm{e}^{-j\frac{\pi u (n+1)(n+2)}{63}}, & 31 \leq n \leq 61 \end{cases} \tag{4.30}$$

where u is called a root index. In LTE, u is used to determine part of the cell identity and can be one of three values. Typically, the receiver does not know which one of the sequences it received, but it can generate local copies of each three and try to cross-correlate with each one. A property of Zadoff-Chu sequences in general is that the absolute value of the cross-correlation between to different sequences can be limited if the roots are chosen according to a specific rule. Consequently, given that the set of possible root indices are chosen wisely, it is possible for the receiver to determine which sequence was received. However, the particular choice of u has additional effects, such as robustness to frequency offsets [77, 78]. Balancing the effects of different set of roots, the LTE standard eventually settled at $u \in \{25, 29, 34\}$.

These are mapped symmetrically around the DC null subcarrier, regardless of the number of total subcarriers. As the total number of subcarriers in use at a given time instant can vary, this allows the receiver to locate the PSS without known how many subcarriers are in total use. Lastly, note that the two halves on either side is a reflection of the other. As we have seen before, these types of structures can be used to find the frequency offset.

From (4.30), it is clear why the has constant amplitude. This is a useful property, as it will bound the amplitude of the time-domain signal and limit the PAPR, easing requirements on amplifiers. In Figure (4.21a), the left part depicts the symbols generated by (4.30). In Figure 4.21b, the same Zadoff-Chu sequence have been transformed into time-domain via the IDFT. In this figure, we see that the first haft of the symbol is mirrored in the second half.

The zero autocorrelation property is shown in to the right in Figure 4.21a. The figure shows the result of (4.5) and we see that indeed, the sequence produces a correlation spectrum which is zero at every other location than the zero lag. In [63], this property is shown analytically and the

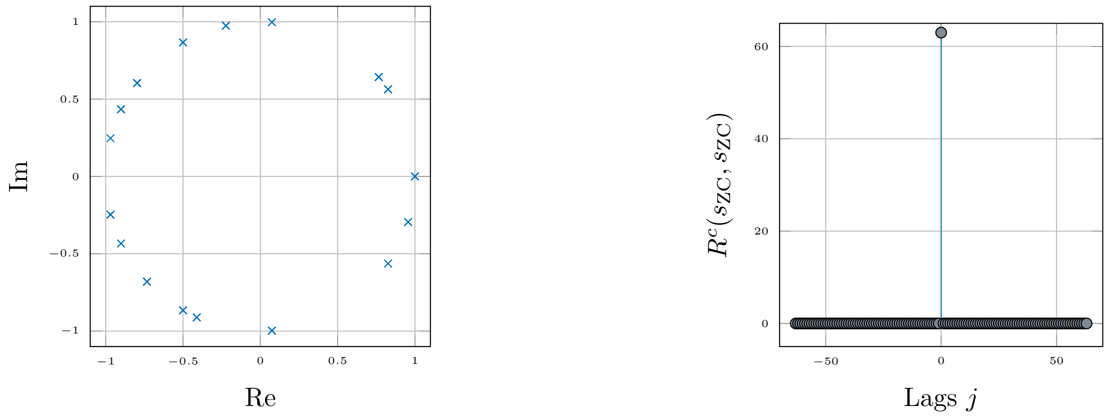

(a) The sequence s_{CZ} in the frequency-domain and its autocorrelation.

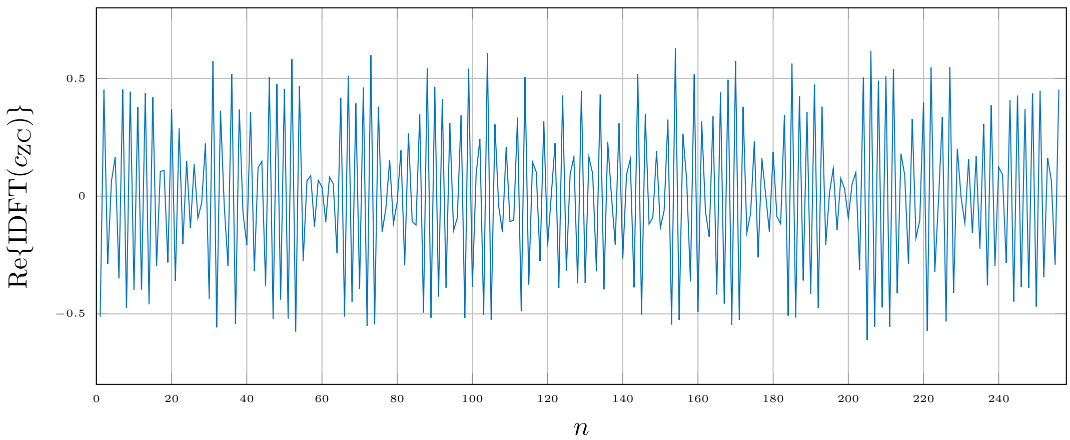

(b) Time-domain representation of s_{CZ}.

Figure 4.21: The autocorrelation property of a Zadoff-Chu sequence. The sequence s_{ZC} is a Zadoff-Chu sequence of length 63 and root $u = 29$ and its properties. The sequence is mapped to the central subcarriers in a system with 256 subcarriers.

interested reader is referred to that paper for details.

In Figure 4.22, we show the cross-correlation between a sequence with root index $u = 29$ and the other two sequences in LTE, as well as with itself.

Note that the correlation is no longer exactly zeros at other lags than $j = 0$, which is due to the fact that we have transformed the sequence into the time-domain. Further, the cross-correlation is not circular, so at every position other than zero lag, we only use $N - |j|$ values in the summation, so the Zadoff-Chu sequence will not completely cancel itself out. Still, we see that the peak at zero lag generated with cross-correlating the sequence with root 29 with itself is significantly higher than the other peaks. Consequently, we will be able to both determine which of the three sequences that is in use, as well as the correct timing.

In Figure 4.23, we show the MSE of the symbol timing estimator. We use the same type of correlation method as in previous cases and we can clearly see an improvement over previous

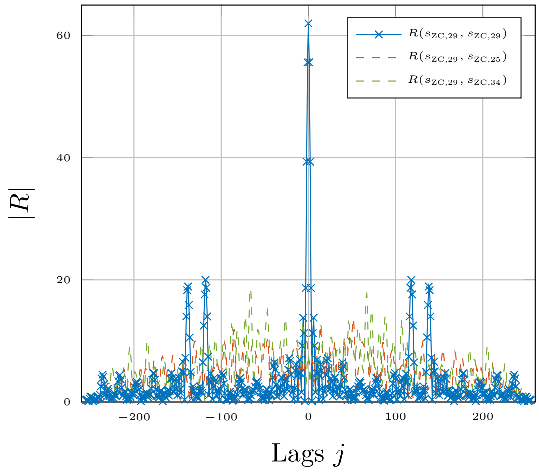

Figure 4.22: The cross-correlation between a Zadoff-Chu sequence with root index $u = 29$ and the two other sequences used in LTE.

results. In fact, at SNR larger than 8 dB in the infinite-precision case and larger than 12 dB in the one-bit case, no errors were recorded at all during the simulations. Comparing this to the previous results, the improvement is striking. It is also interesting to note that the quantized case also gains significantly compared to the earlier results. This is highly promising, as it hints towards the viability of using one-bit quantizers in a real-world system.

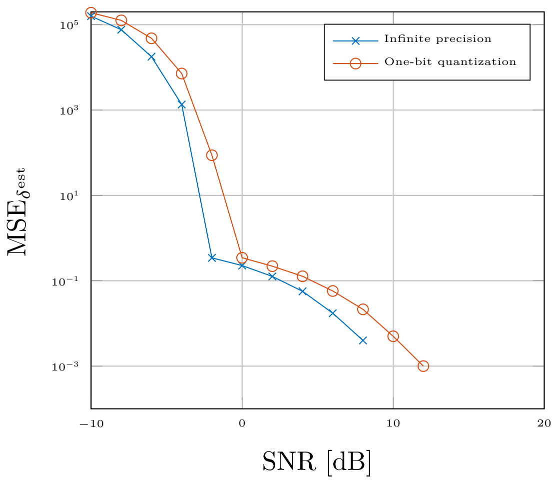

Figure 4.23: The MSE performance of the STO estimator using a Zadoff-Chu sequence. Results obtained over the AWGN channel with 256 subcarriers.

Next, we examine the robustness of the method. Similar to the previous method, we again gather the peak value of the correlation spectrum when the signal is present and when it is not. We construct the histogram and the cumulative histogram, as shown in Figure 4.24.

In Figure 4.24a, we see an approximation of the distribution of the maximum value of the correlation spectrum at both the correct and incorrect timing. Note the scaling of the x-axis however, as it from a quick visual inspection might appear as if the result is similar in both cases. In the infinite-precision case, the distribution is more separated than in the one-bit case, echoing a

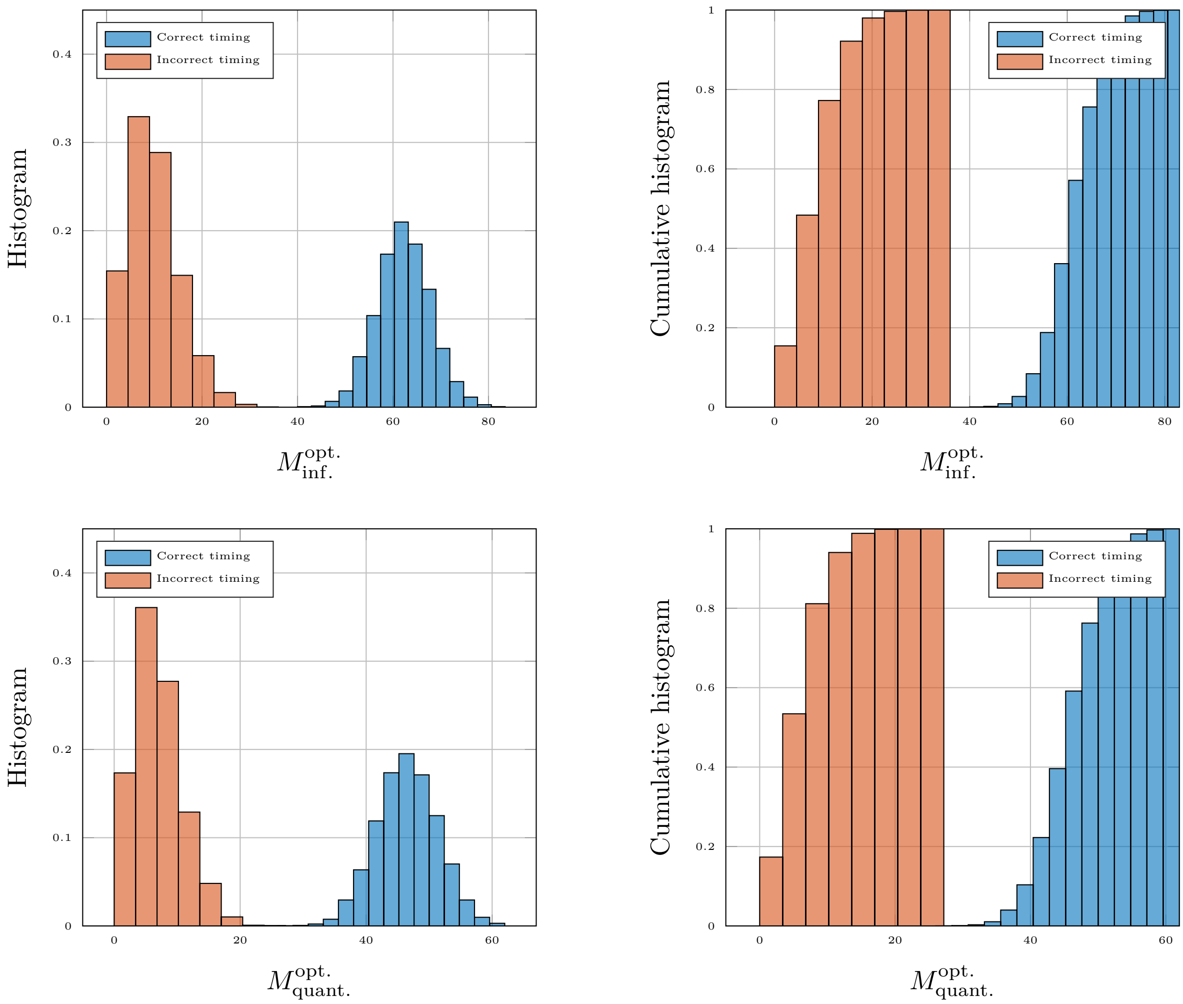

(a) Normalized histograms. **(b)** Normalized cumulative histograms.

Figure 4.24: Histogram and cumulative histogram of infinite-precision estimator (top) and one-bit estimator (bottom). Results obtained via recording $5 \cdot 10^4$ values of each estimator when there was a signal and when there was not. SNR set to 0 dB.

familiar result. In Figure 4.24b, the cumulative histograms are depicted and as there is no overlap in neither this figure or Figure 4.24a, the risk of false detection seems small.

In Figure 4.25, the missed probability and false detection probability is shown. Clearly, a threshold can be chosen so that the risk of missing a valid sequence while still rejecting anything that is not the synchronization signal is almost arbitrarily small.

Up until this point, we have only focused on the PSS. This is because this symbol is what it used to achieve the initial timing estimation. The SSS is then used to for example deduce the full cell identity, as well as a number of other things. The second symbol is constructed of two known and interlaced m-sequences, generating a symbol with PN characteristics. Using these kinds of symbols for synchronization have been investigated in for example [79], where it was shown that preambles with PN-sequences can outperform training symbols with repeated parts, even in the one-bit case.

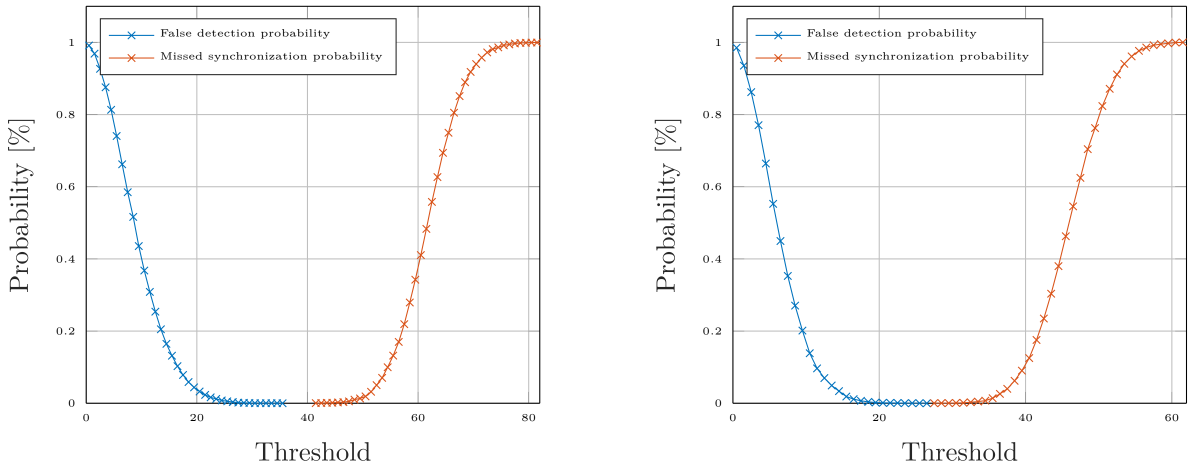

(a) Infinite-precision case.

(b) One-bit quantization case.

Figure 4.25: The normalized histogram and cumulative histograms in the infinite-resolution and one-bit quantization case. Results obtained via recording $5 \cdot 10^4$ values of each estimator when there was a signal and when there was not. SNR set to 0 dB.

Exactly how to use to PSS and SSS to achieve synchronization is, as mentioned in the beginning of this section, not given by the standard. It is interesting to look at the correlation properties of the PSS in the one-bit case, as this most definitely must be used to achieve the initial synchronization. For the other parts, a number of papers have proposed different schemes, so it can perhaps be assumed that some variant of these ideas is present in current systems. In [80], a scheme where the CP is used to obtain an estimate of the STO and the fractional CFO by maximizing (4.7). The SSS is then used to determine the integer CFO in a manner similar to Schmidl and Cox. In [56], the CP is used to obtain rough estimates of the STO and CFO which are later refined via the PSS and SSS. Here, the fact that the PSS consists of the mirrored halves are used to find the frequency offset. Another method was suggested in [81], where both the PSS and SSS are used together to achieve carrier frequency synchronization. In [82], a ML approach to estimate several parameters is demonstrated. This method outperforms a number of other, more heuristic methods, at the expense of computational complexity.

We have now investigated a number of different methods and given an overview of the topic itself. For convenience, we here give Figure 4.26, where the MSE for the STO estimation for the methods considered is shown in Figure 4.26a, and the MSE for the CFO estimators in Figure 4.26b. Comparing the CP-based method and the Schmidl and Cox method, we see that their MSE performance is similar, with a slight edge for the CP-based method. The main advantages of the Schmidl and Cox method, however, is the robustness, which is not shown here.

Further, we note that the LTE-like method with a Zadoff-Chu sequence has vastly better MSE behavior than the others for the STO estimator. This demonstrates that the choice of synchronization sequence greatly influences the performance. Also, using a single known symbol instead of a symbol with repetitive structure, we get a single peak in the correlation spectrum rather than a plateau. This allows the MSE to become arbitrarily low at high SNRs. We did not do any frequency estimation with the PSS and SSS, but note that the framework available via the CP, Zadoff-Chu and m-sequence is enough to construct a method matching the performance of

Morelli and Mengali. Fundamentally, Morelli and Mengali were able to improve the accuracy by looking at symbol with more than one repetition, so that a better estimate may be produced via a weighted average. With both the PSS and SSS available, similar ideas can be pursued and in the reference list, some successful examples are given.

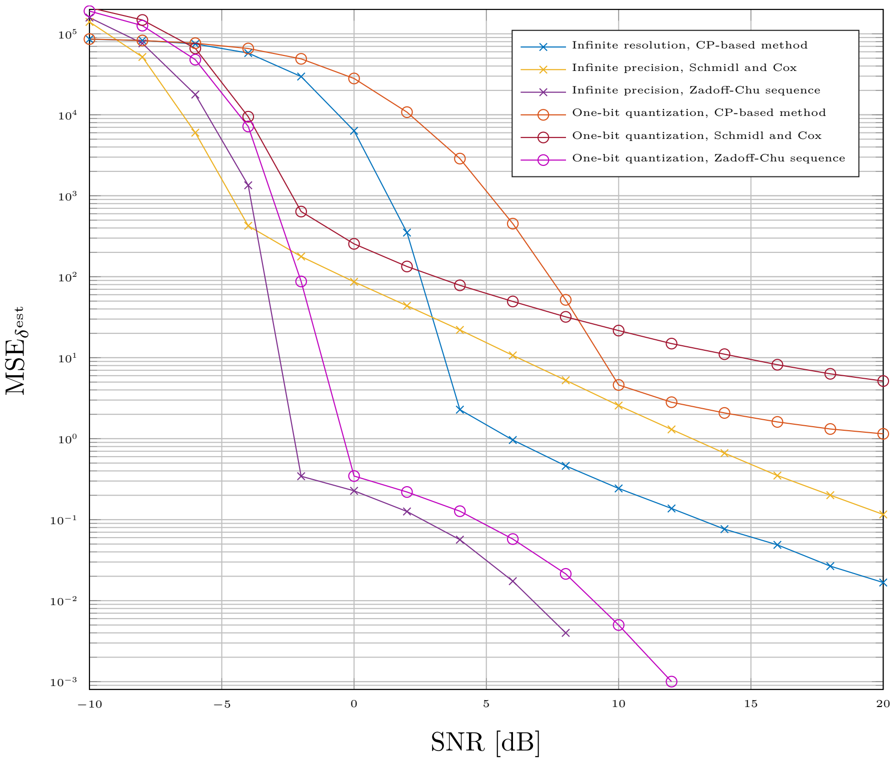

(a) The MSE of the timing estimators.

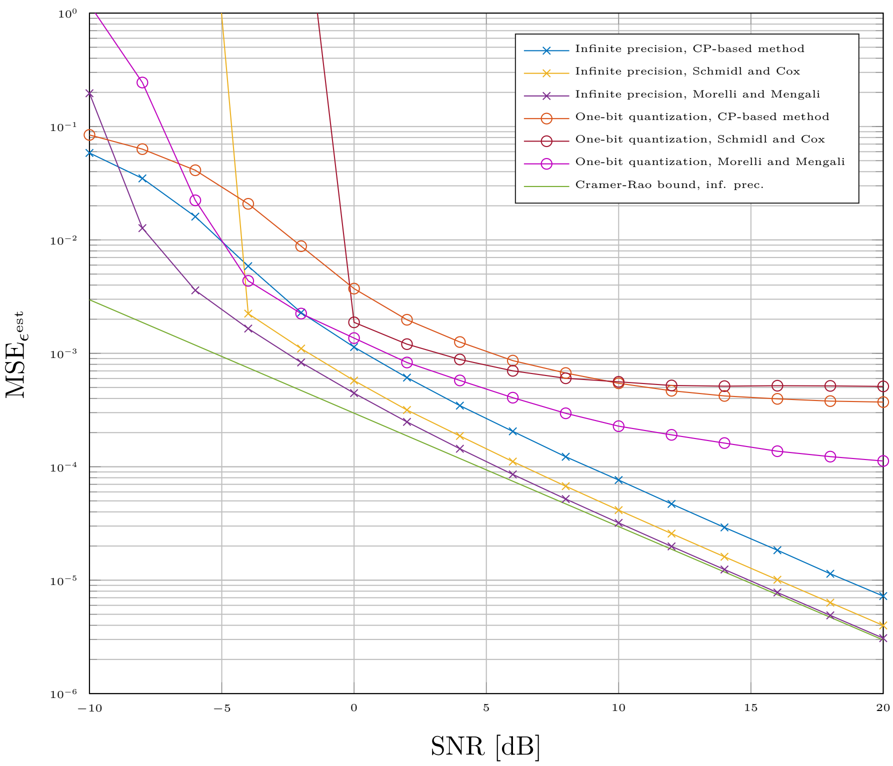

(b) The MSE of the frequency estimators, $\epsilon_0 = 0.2$.

Figure 4.26: MSE for the different STO and CFO estimators. Results obtained from simulation over an AWGN channel, 512 subcarriers, cyclic prefix of length 36, Zadoff-Chu sequence of length 63 and $5 \cdot 10^4$ trials at each SNR.

5

Performance

In this chapter, we will examine the performance of the system under STO and CFO. A MU-MIMO system will be simulated and the the resulting SINDR for a user will be compared with the analytic expression for the SINDR found in Chapter 3. Lastly, we will look at the root mean square error (RMSE) of the STO and CFO estimation.

We begin with the received SINDR when there is STO present. In Figure 5.1, we see the results for both the infinite-precision case and the one-bit case both when there is CFO and when there is not. Here, we assume that there is residual STO after the estimation, denoted by $\delta_u = \delta_{0,u} - \delta_u^{\text{est}}$.

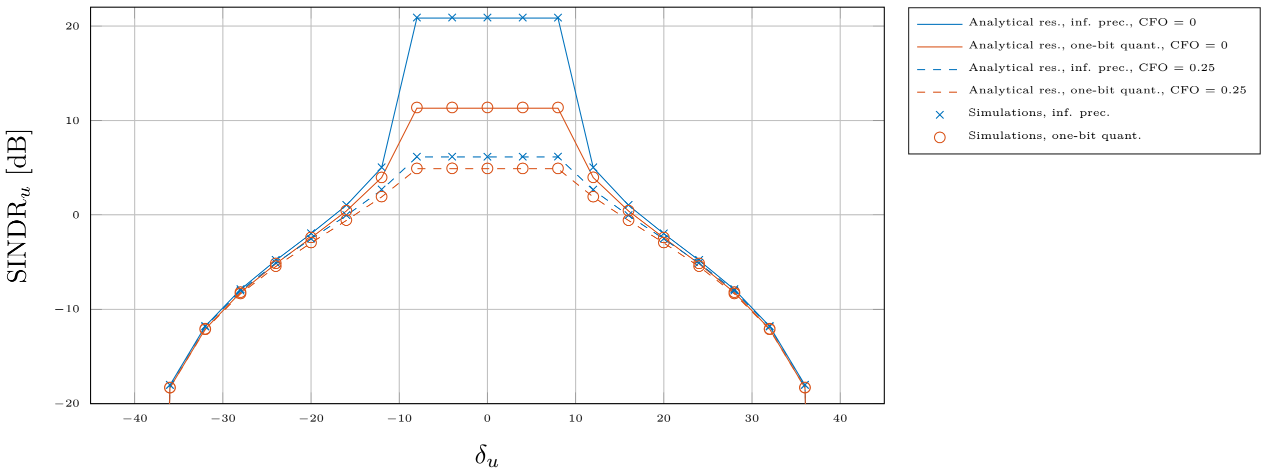

Figure 5.1: The effect on the SNR under STO from analysis and simulations. The solid and dashed lines are analytic results and the marks are simulations. Results obtained for MU-MIMO system with a flat-fading channel model, 128 BS antennas, 8 UEs and 32 subcarriers, cyclic prefix of length 16 and SNR at 0 dB.

As predicted by the analysis, the received SINDR quickly decreases as the STO estimation error grows. Because of the asymmetrical effect of the STO, it is commonplace to slightly shift the starting point of the DFT window. If the start of the window is taken to be somewhere in the middle of the CP, or perhaps slightly tilted toward the end of the CP in a frequency-selective channel, the received SINDR will be unaffected by small STO estimation errors in either direction. As there is no reason to assume that the estimation error should be tilted toward the negative side, this yields a more robust system. This has been done in here, explaining why Figure 5.1 is symmetrical around $\delta_u = 0$. Note that there is no reduction in the received SINDR as long as the

STO causes the start of the DFT window to be within the CP.

Moreover, in Figure 5.1, we see that the one-bit case has an approximate 10 dB loss compared to infinite-precision case in the case of no CFO. This is due to the quantization distortion, which as we have seen throughout this thesis limits the performance in a quantized case. It is however interesting to note that the effect of STO is similar, regardless of the resolution of the quantization. Additionally, the effect of STO is in this case more severe than quantization distortion, highlighting the importance of accurate synchronization. Lastly, adding CFO to the system, we note a significant drop in the performance, again underscoring the need for accurate synchronization.

Next, we examine the effect of CFO in the system. In Figure 5.2, results are shown. This effect is completely symmetrical, as rotations in a particular direction is not worse than the other. As noted in previous chapters, the fractional CFO is what causes ICI, so the worst case scenario for the residual CFO $\epsilon_u = |\epsilon_{0,u} - \epsilon_u^{\text{est}}|$ is 1/2 subcarrier spacings.

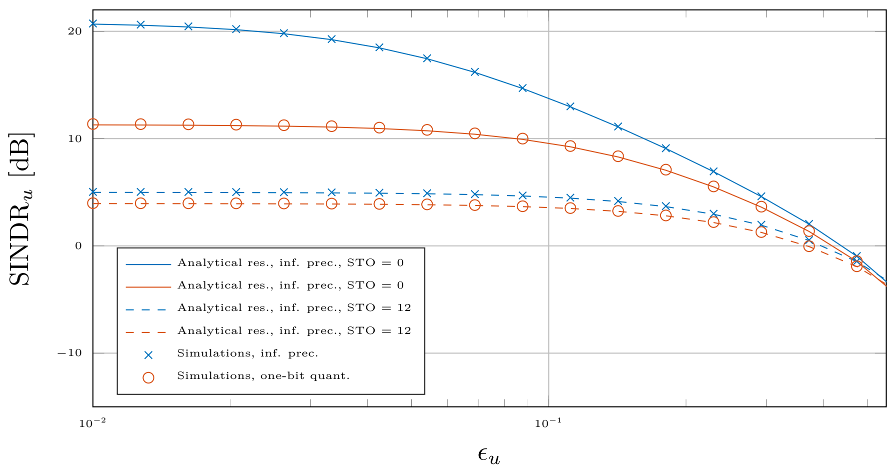

Figure 5.2: The effect on the SNR under CFO from analysis and simulations. The solid and dashed lines are analytic results and the marks are simulations. Results obtained for MU-MIMO system with a flat-fading channel model, 128 BS antennas, 8 UEs and 32 subcarriers, cyclic prefix of length 16 and SNR at 0 dB.

Here, we see for example that a CFO of 0.15 subcarrier spacing causes a loss of about 10 dB in the infinite-precision case. This illustrates the inherent sensitivity to CFO in the OFDM scheme mentioned in Chapter 2 and why significant effort has been put into the design of effective CFO estimation strategies.

Moreover, we again see an additional penalty due to the quantization distortion of about 10 dB. In the figure, we note that the effect of the CFO does not seem to be as severe as in the infinite-precision case, at least in a relative sense. Looking at the case when the residual CFO is 0.2, we see that the SINDR in the quantized is approximately 4 dB. Comparing this to the around14 dB loss in the infinite-precision case, it is interesting to note that since the one-bit case suffers from a significant distortion due to due quantization itself, the effect of the CFO is less noticeable. This is similar to what we saw in the Figure 5.1. In light of these results, we note that the requirements on synchronization in a coarsely quantized system could possibly be more lenient than with fine quantization. As so much distortion is added already, the effect of slight STO or CFO is not very significant.

In Figure 5.3, we see that effect of adding more receiving antennas to the system. The main take-away of Figure 5.3a and Figure 5.3b is that we do get an array gain in both the the infinite-resolution and the quantized case, meaning that more antennas improve the received SNR when there is STO in both the infinite-precision and one-bit case.

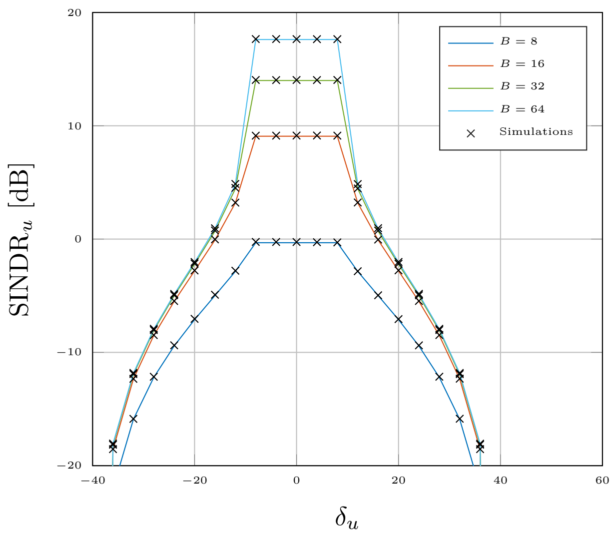

(a) Infinite-precision.

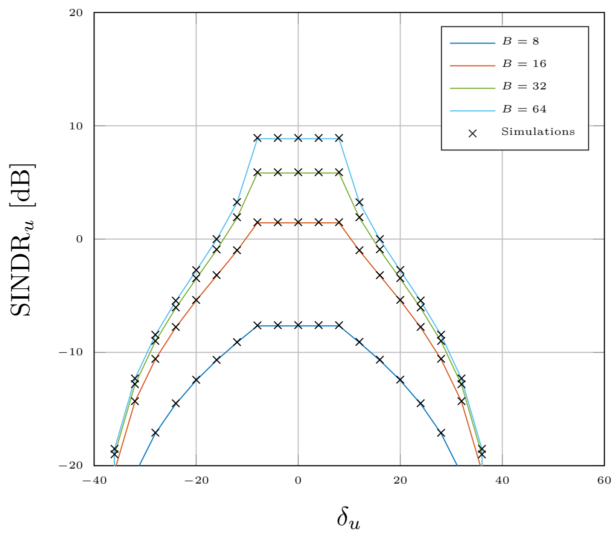

(b) One-bit quantization.

Figure 5.3: The effect on the received SNR when adding more antennas. Results obtained for MU-MIMO system with a flat-fading channel model, CFO set tot zero, 8 UEs and 32 subcarriers, cyclic prefix of length 16, SNR at 0 dB and a varying number of BS antennas B.

Examining Figure 5.3b in particular, we note that the gain of adding additional antennas seems to diminish as the number of antennas B grows. This is due to the fact that while more antennas can average out the effect of thermal noise, it cannot decrease the quantization distortion. This is an interesting phenomena, explained by the fact that if there is correlation between the input to a group of quantizers, the gain of combining their output will be limited. The number of antennas beyond which the received SNR after combing will increase only marginally will depend on the specifics of the channel.

Lastly, we examine the RMSE performance for this system. In Figure 5.4, the results are shown. Since the STO and CFO are estimated per antenna, there is no benefit from adding antennas in the RMSE sense. As noted in Chapter 4, a significant gain in performance can be obtained via some joint processing, where the received signal from each antenna is combined in some way before the offsets are estimated. This would require the offsets to be the same on each antenna, which, at least in certain cases, is probably quite reasonable. That approach was not pursued here and it is noted that the gap between the performance in the infinite-precision and the one-bit case can probably be reduced via more complex processing.

Examining the results in Figure 5.4a and 5.4b, we note a floor in the performance for the quantized case. This is due to the quantization distortion and echos what was found in Chapter 4.

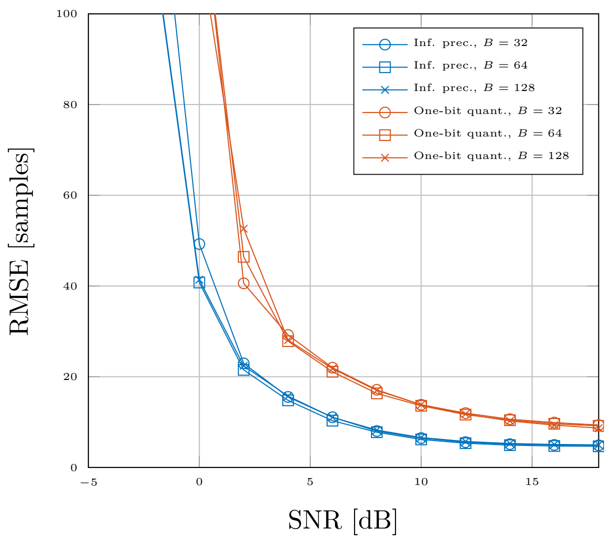

(a) STO

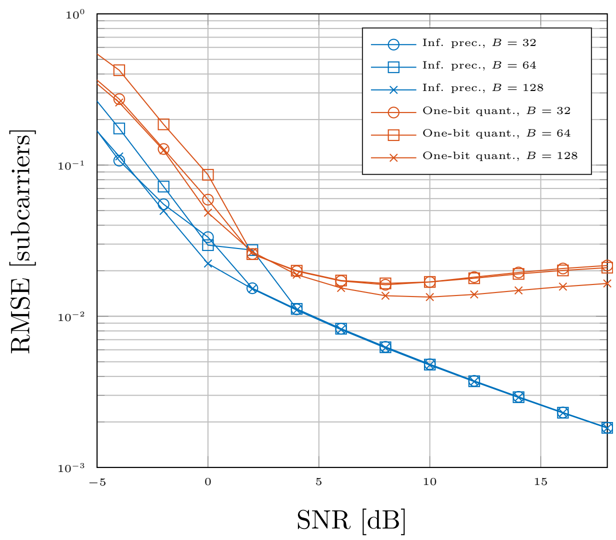

(b) CFO

Figure 5.4: The RMSE for the Schmidl and Cox algorithm in a SIMO system with frequency-selective fading channel model with 10 taps, 1200 occupied subcarriers out of a total of 2048 subcarriers, cyclic prefix of length 144 samples, 10 pilot symbols with LS channel estimation and ZF combining.

It is interesting to note that in Figure 5.4a, the RMSE of the STO estimation in the quantized case is only a few samples worse than the infinite-precision case for the SNR range considered. However, as seen in Figure 5.1, even a small STO can have a significant impact on performance. Perhaps, the cost of a few extra bits of quantization resolution is a reasonable expense to increase the system robustness.

6

Discussion

In this chapter, some reflection on the results obtained throughout this thesis will be presented. A fair amount of insights has already been raised a numerous places, so this chapter intends to summarize, raise some points that have not previously been brought up and mention some concluding remarks. Lastly, we will provide some ideas for future research.

In this thesis, we started by examining the effects of synchronization errors on an OFDM signal if we include a one-bit quantizer in the system. The types of errors considered are symbol timing offset, relating the uncertainty in the exact arrival of a signal, as well as frequency offset, arising from uncertainty regarding the carrier frequency of the signal. Beginning with an overview of the effects, it was noted that both the STO and CFO potentially added significant distortion to our system. In order to be able to properly evaluate the simulations results, the next step was to derive the analytic expressions for the received signal under these types of errors. In the literature review stage of this thesis, no example of an expression involving both the STO and CFO was found – the effects were always treated separately. Consequently, to find a more general expression was the starting point of the analysis. In Chapter 3, the analysis was carried out and an expression describing both these effects in i MU-MIMO scenario was found. It was showed that the findings agree with previous results, by setting one of the parameters to zero and note that the analysis revealed an interaction between the STO and CFO that, to the author's knowledge, had not been previously documented. Moreover, as the analysis was limited to a flat-fading channel model and symbol-sampling rate, the extension to the one-bit case via Bussgang's Theorem was straightforward and a theoretical expression for the SINDR in the presence of STO and CFO was found.

In the following chapter, the topic of synchronization was examined. It was noted that the optimum method for detecting a synchronization sequence reduced to a simple correlation in the one-bit quantization case. Next, several synchronization methods were reviewed and the main conclusion from this chapter was that none of the methods failed when introducing the one-bit quantization. In fact, the quantized case behaved similarly in most cases, demonstrating that the effect of quantization in the context reduces to decreasing the SINDR. However, it was detected that while most STO estimation methods are insensitive to the CFO in the infinite precision case, this is not true in the quantized case. The level of CFO strongly influenced the performance of the estimators in the one-bit case and an intuitive explanation for the effect was presented. This is interesting, as it is common for STO and CFO estimation method to assume that these can be done separately. In the one-bit case, taking the fact that the performance of the described

STO estimation methods are influenced by the level of CFO into account may yield significant performance gains. One idea might be to have an algorithm which estimates first the STO and CFO and then after correcting the received signal according to the estimations, estimates the offset again. This algorithm could repeat until the difference between two successive estimations is sufficiently small or some maximum number of iterations have been met.

The main takeaway from Chapter 4 was that standard methods, such as using the cyclic prefix, a repeated symbol or a dedicated known preamble, does seem to work in one-bit quantized cases. This is a significant finding and, combined with other result from previous research, indicates that low-resolution quantizers is indeed feasible way to mitigate the power consumption issues of large antenna arrays. It is already specified that some parts of future 5G systems will run on frequencies than higher than 6 GHz and consequently, it seems all but certain that large antenna arrays will be included in future communication system. To facilitate the beamforming necessary for acceptable SNR levels, a large number of antennas will be required. To keep the energy bill at a reasonable level, the energy consumption per part in the antenna hardware chain must be reduced. As the quantizer is responsible for a sizable share of the needed energy, lowering either its speed or resolution is an attractive option. It seems unlikely that low-speed quantizers will ever yield acceptable performance, so low-resolution quantizers is probably the most likely path forward. The results of this thesis, which indicate that that systems employing these types of quantizers can indeed be synchronized, speaks to the feasibility of this solution.

Lastly, the performance of systems with either infinite resolution or one-bit quantization was examined. It was noted that the SNR drop in the one-bit case was relatively speaking significantly lower than in the infinite precision case. This is interesting and means that as the quantization already adds a large amount of distortion, the added interference from STO or CFO is marginal in relation the the distortion. This would mean that perhaps, the requirements on synchronization can be less strict in the one-bit case than in the infinite precision case. This in turn would possibly mean that the CP can be made slightly shorter, thereby reducing the overhead and increasing the throughput.

6.1 Ideas for future research

When considering future research, a number of ideas spring to mind. For example, while deriving the SINDR in the presence of simultaneous STO and CFO, some limiting assumptions were made. In particular, there are some areas where there are fairly obvious extensions to the current analysis that were not pursued due to time considerations.

Firstly, in this thesis, the channel was assumed to be flat-fading. The extension to a more general frequency-fading channel is however quite straightforward and requires basically the same steps as shown in Chapter 3. Generalizing the analysis to this case and examining the performance under these conditions is an obvious continuation of this work and an important step to increase the understanding of this topic.

Another limitation in this thesis is the assumption of symbol-sampling rate. In several papers, e.g. [21], the benefits of oversampling in the one-bit quantization case has been demonstrated. Consequently, it is of great interest to examine if the gains of oversampling are present when synchronization errors are introduced to the system. Including oversampling would mean that $\mathbb{E}\left[\hat{x}[k]\hat{x}[k']'\right]\right] \neq 0, k \neq k'$, complicating the analysis.

Next, in the area of synchronization, there is some analysis missing in the one-bit case. This could be useful to validate the results of the simulations, as well as the further deepen the understanding of the effect of one-bit quantization. Specifically, the development of a maximum likelihood method relying on the cyclic prefix, analogous to [53] would be a highly interesting addition to the current knowledge base.

The main point to mention in this section, however, is that the time has come for the examination of the synchronization methods in a real-world system. While the results of this thesis strongly hints towards the fact that standard methods, such as Schmidl and Cox, will work in one-bit system, no definitive answer can be given until real-world experiments have been carried out. For example, this thesis have not examined the effects of sampling clock offsets, which is a well-known issue in many real system. How much this affects the system is yet unknown and perhaps this is a prerequisite for real-world experiments. The initial intention of this thesis was to begin the experimentation phase, but the decision to exclude that from the scope had to be taken with time constraints in mind.

It is the final conclusion of this thesis that standard methods can be used in systems employing one-bit quantization. An initial examination of this topic have been carried out and some important findings have been highlighted, such as that the quantization distortion causes so much distortion that the additional effect of a small STO or CFO is insignificant. We also note the CFO dependence of the performance of the timing estimator, as well as an analytic expression for the SINDR under STO and CFO.

Bibliography

[1] Nationalencyklopedin. (2018) 2G. [Online]. Available: https://www.ne.se/uppslagsverk/encyklopedi/lång/2g

[2] NTT DoCoMo. (2004) DoCoMo overview. [Online]. Available: https://www.oecd.org/sti/ieconomy/32145299.pdf

[3] Ericsson. (2018) Ericsson 5G. [Online]. Available: https://www.ericsson.com/en/5g

[4] A. Osseiran, F. Boccardi, V. Braun, K. Kusume, P. Marsch, M. Maternia, O. Queseth, M. Schellmann, H. Schotten, H. Taoka, H. Tullberg, M. A. Uusitalo, B. Timus, and M. Fallgren, "Scenarios for 5G mobile and wireless communications: the vision of the METIS project," *IEEE Commun. Mag.*, vol. 52, no. 5, pp. 26–35, May 2014.

[5] J. G. Andrews, S. Buzzi, W. Choi, S. V. Hanly, A. Lozano, A. C. K. Soong, and J. C. Zhang, "What will 5G be?" *IEEE J. Sel. Areas Commun.*, vol. 32, no. 6, pp. 1065–1082, Jun. 2014.

[6] Reuters. (2018, feb) Olympics - the first 5G Olympics? Not quite, say bemused spectators. [Online]. Available: https://www.reuters.com/article/us-olympics-2018-5g/olympics-the-first-5g-olympics-not-quite-say-bemused-spectators-idUSKCN1G80C3

[7] Ericsson. (2018) The 5G race is on. [Online]. Available: https://www.ericsson.com/en/news/2017/5/ericsson-and-verizon-test-the-limits-of-5g

[8] F. Boccardi, R. W. Heath Jr., A. Lozano, T. L. Marzetta, and P. Popovski, "Five disruptive technology directions for 5G," *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 74–80, Feb. 2014.

[9] S. Arnon, Ed., *Visible Light Communication.* Cambridge University Press, 2015.

[10] *3GPP TS 38.101-1 Release 15*, 3GPP Std., Sep. 2018.

[11] *3GPP TS 38.101-2 Release 15*, 3GPP Std., Sep. 2018.

[12] H. Friis, "A note on a simple transmission formula," *Proc. IRE*, vol. 34, no. 5, pp. 254–256, May 1946.

[13] Q. Li, G. Li, W. Lee, M. il Lee, D. Mazzarese, B. Clerckx, and Z. Li, “MIMO techniques in WiMAX and LTE: a feature overview,” *IEEE Commun. Mag.*, vol. 48, no. 5, pp. 86–92, May 2010.

[14] F. Rusek, D. Persson, B. K. Lau, E. G. Larsson, T. L. Marzetta, O. Edfors, and F. Tufvesson, “Scaling up MIMO,” *IEEE Signal Process. Mag.*, vol. 30, no. 1, pp. 40–60, Jan. 2013.

[15] E. G. Larsson, O. Edfors, F. Tufvesson, and T. L. Marzetta, “Massive MIMO for next generation wireless systems,” *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 186–195, Feb. 2014.

[16] S. Jacobsson, G. Durisi, M. Coldrey, U. Gustavsson, and C. Studer, “Throughput analysis of massive MIMO uplink with low-resolution ADCs,” *IEEE Trans. Wireless Commun.*, vol. 16, no. 6, pp. 4038–4051, Apr. 2017.

[17] R. H. Walden, “Analog-to-digital converter survey and analysis,” *IEEE J. Sel. Areas Commun.*, vol. 17, no. 4, pp. 539–550, Apr. 1999.

[18] ——, “Analog-to-digital converter technology comparison,” in *Proc. IEEE GaAs IC Symp.* Phildelphia, PA, USA: IEEE, Aug. 1994, pp. 217–219.

[19] B. Murmann. (2018) ADC performance survey 1997-2018. [Online]. Available: http://web.stanford.edu/~murmann/adcsurvey.html

[20] T. Instruments. (2016) High-speed ADCs (>10MSPS) – RF sampling. [Online]. Available: http://www.ti.com/data-converters/adc-circuit/high-speed/rf-sampling.html

[21] S. Jacobsson, G. Durisi, M. Coldrey, and C. Studer, “Massive MU-MIMO-OFDM downlink with one-bit DACs and linear precoding,” in *IEEE Global Commun. Conf. (GLOBECOM).* Singapore, Singapore: IEEE, Dec. 2017.

[22] C. Mollén, “High-end performance with low-end hardware,” Ph.D. dissertation, Linköping University, 2017.

[23] M. Sarajlic, L. Liu, and O. Edfors, “When are low resolution ADCs energy efficient in massive MIMO?” *IEEE Access*, vol. 5, pp. 14 837–14 853, Jun. 2017.

[24] A. Wadhwa and U. Madhow, “Blind phase/frequency synchronization with low-precision ADC: A Bayesian approach,” in *Allerton Conf. Commun., Contr., Comp.* Monticello, IL, USA: IEEE, Oct. 2013.

[25] D. Zhu and R. W. Heath Jr, “Directional timing synchronization in wideband millimeter wave cellular systems with low-resolution ADCs,” in *Asilomar Conf. on Sig., Sys., and Comp.* Pacific Grove, CA, USA: IEEE, Oct. 2017.

[26] D. Zhu, R. Bendlin, S. Akoum, A. Ghosh, and R. W. Heath Jr, “Double-sequence frequency synchronization for wideband millimeter-wave systems with few-bit ADCs,” *ArXiv e-prints*,

Dec. 2018.

[27] H. Wang, W.-T. Shih, C.-K. Wen, and S. Jin, “Reliable OFDM receiver with ultra-low resolution ADC,” *IEEE Transactions on Communications*, pp. 1–1, Jan. 2019.

[28] F. Tomatis and S. Sesia, *LTE - The UMTS Long Term Evolution*, 2nd ed. Wiley John and Sons, 2011, ch. Synchronization and Cell Search, pp. 151–163.

[29] R. W. Chang, “Synthesis of band-limited orthogonal signals for multichannel data transmission,” The Bell Sys. Tech. Jour., Dec. 1966.

[30] S. B. Weinstein, “The history of orthogonal frequency-division multiplexing,” *IEEE Communications Magazine*, vol. 47, no. 11, pp. 26–35, nov 2009.

[31] J. A. C. Bingham, “Multicarrier modulation for data transmission: An idea whose time has come,” *IEEE Commun. Mag.*, vol. 28, no. 5, pp. 5–14, May 1990.

[32] Y. S. Cho, J. Kim, W. Y. Yang, and C. G. Kang, *MIMO-OFDM Wirelss Communications With MATLAB*, Wiley, Ed. John Wiley & Sons Inc., 2010.

[33] C. D. Perels, “Frame-based MIMO-OFDM systems: Impairment estimation and compensation,” Ph.D. dissertation, ETH Zurich, 2007.

[34] S. Jacobsson, G. Durisi, M. Coldrey, U. Gustavsson, and C. Studer, “One-bit massive MIMO: Channel estimation and high-order modulations,” in *IEEE Int. Conf. Commun. Works. (ICCW)*. London, United Kingdom: IEEE, Jun. 2015, pp. 1304–1309.

[35] R. W. Heath Jr., N. González-Prelcic, S. Rangan, W. Roh, and A. M. Sayeed, “An overview of signal processing techniques for Millimeter Wave MIMO systems,” *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 3, pp. 436–453, Apr. 2016.

[36] S. Krone and G. Fettweis, “Achievable rate with 1-bit quantization and oversampling at the receiver,” in *IEEE Commun. Theo. Worksh.*, Cancun, Mexico, May 2010.

[37] Y. Li, C. Tao, G. Seco-Granados, A. Mezghani, A. L. Swindlehurst, and L. Liu, “Channel estimation and performance analysis of one-bit Massive MIMO systems,” *IEEE Trans. Signal Process.*, vol. 65, no. 15, pp. 4075–4089, Aug. 2017.

[38] J. J. Bussgang, “Crosscorrelation functions of amplitude-distorted Gaussian signals,” MIT Res. Lab. of Elec., Cambridge, MA, Mar. 1952.

[39] J. Singh, O. Dabeer, and U. Madhow, “Communication limits with low precision analog-to-digital conversion at the receiver,” in *IEEE Int. Conf. Commun. (ICC)*. Glasgow, United Kingdom: IEEE, jun 2007.

[40] S. Jacobsson, G. Durisi, M. Coldrey, T. Goldstein, and C. Studer, “Quantized precoding for

massive MU-MIMO," *IEEE Trans. Commun.*, vol. 65, no. 11, pp. 4670–4684, Nov. 2017.

[41] H. E. Rowe, "Memoryless nonlinearities with Gaussian inputs: Elementary results," The Bell Sys. Tech. Jour., Sep. 1982.

[42] Y. Mostofi and D. C. Cox, "Mathematical analysis of the impact of timing synchronization errors on the performance of an OFDM system," *IEEE Trans. Commun.*, vol. 54, no. 2, pp. 226–230, Feb. 2006.

[43] J. H. V. Vleck and D. Middelton, "The spectrum of clipped noise," *Proc. IEEE*, vol. 54, no. 1, pp. 2–19, Jan. 1966.

[44] T. Pollet, M. V. Bladel, and M. Moeneclaey, "BER sensitivity of OFDM systems to carrier frequency offset and Wiener phase noise," *IEEE Trans. Commun.*, vol. 43, no. 2/3/4, pp. 191–193, Feb. 1995.

[45] W. D. Warner and C. Leung, "OFDM/FM frame synchronization for mobile radio data communication," *IEEE Trans. Veh. Technol.*, vol. 42, no. 3, pp. 302–313, Aug. 1993.

[46] P. H. Moose, "A technique for orthogonal frequency division multiplexing frequency offset correction," *IEEE Trans. Commun.*, vol. 42, no. 10, pp. 2908–2914, Oct. 1994.

[47] H. Nogami and T. Nagashima, "A frequency and timing period acquisition technique for OFDM systems," in *Proc. Pers., Ind. Mob. Rad. Commun. (PIMRC)*, vol. 3. IEEE, Sep. 1995, pp. 1010–1015.

[48] T. M. Schmidl and D. C. Cox, "Robust Frequency and Timing Synchronization for OFDM," *IEEE Trans. Commun.*, vol. 45, no. 12, pp. 1613–1621, Dec. 1997.

[49] ——, "Low-overhead, low-complexity [burst] synchronization for OFDM," in *Proc. IEEE Int. Conf. on Commun. (ICC/SUPERCOMM)*, vol. 3. Dallas, TX, USA: IEEE, Jun. 1996, pp. 1301–1306.

[50] K. Wang, M. Faulkner, and J. Singh, "Timing synchronization for 802.11a WLANs under multipath channels," in *Proc. Australian Telecomm. Netw. and Appl. (ATNAC)*, Melbourne, Australia, Dec. 2003.

[51] M. Morelli and U. Mengali, "An improved frequency offset estimator for OFDM applications," *IEEE Commun. Lett.*, vol. 3, no. 3, pp. 75–77, Mar. 1999.

[52] M. Sandell, J.-J. van de Beek, and P. O. Börjesson, "Timing and frequency synchronization in OFDM systems using the cyclic prefix," in *Proc. IEEE Int. Symp. on Synch.*, Essen, Germany, Dec. 1995, pp. 16–19.

[53] J.-J. van de Beek, M. Sandell, and P. O. Börjesson, "ML estimation of time and frequency offset in OFDM systems," *IEEE Trans. Signal Process.*, vol. 45, no. 7, pp. 1800–1805, Jul.

1997.

[54] J. L. Massey, "Optimum frame synchronization," *IEEE Trans. Commun.*, vol. 20, no. 2, pp. 115–119, Apr. 1972.

[55] P. T. Nielsen, "Some optimum and suboptimum frame synchronizers for binary data in Gaussian noise," *IEEE Trans. Commun.*, vol. 21, no. 6, pp. 770–772, Jun. 1973.

[56] W. Xu and K. Manolakis, "Robust synchronization for 3GPP LTE system," in *IEEE Glob. Telecomm. Conf. (GLOBECOM).* Miami, FL, USA: IEEE, Dec. 2010.

[57] G. L. Lui and H. H. Tan, "Frame synchronization for Gaussian channels," *IEEE Trans. Commun.*, vol. 35, no. 8, pp. 818–829, Aug. 1987.

[58] B. H. Moon and S. S. Soliman, "ML frame synchronization for the Gaussian channel with ISI," in *Proc. IEEE Int. Conf. Commun. (ICC).* Denver, CO, USA: IEEE, Jun. 1991.

[59] P. Robertson, "Maximum likelihood frame synchronization for flat fading channels," in *Proc. IEEE Int. Conf. on Commun. (ICC/SUPERCOMM).* Chicago, IL, USA: IEEE, Jun. 1992.

[60] M. Speth, F. Classen, and H. Meyr, "Frame synchronization of OFDM systems in frequency selective fading channels," in *Proc. IEEE Veh. Tech. Conf. (VTC)*, vol. 3. Phoenix, AZ, USA,: IEEE, May 1997, pp. 1807–1811.

[61] R. Scholtz, "Frame synchronization techniques," *IEEE Trans. Commun.*, vol. 28, no. 8, pp. 1204–1213, Aug. 1980.

[62] J. Jedwab, "What can be used instead of a Barker sequence?" *Contemporary Mathematics*, vol. 461, pp. 153–178, Feb. 2008.

[63] D. C. Chu, "Polyphase codes with good periodic correlation properties," *IEEE Trans. Inf. Theory*, vol. 18, no. 4, pp. 531–532, Jul. 1972.

[64] K. H. Leung and B. Schmidt, "The field descent method," *Designs, Codes and Cryptography*, vol. 36, no. 2, pp. 171–188, Aug. 2005.

[65] D. Terr and E. W. Weisstein. (2004) Barker code. [Online]. Available: http://mathworld.wolfram.com/BarkerCode.html

[66] T. Helleseth, "50 years of crosscorrelation of m-sequences," in *Int. Castle Meet. Cod. Theo. and Appl. (ICMCTA)*, Estonia, Aug. 2017.

[67] ——, "Linear and nonlinear sequences and applications to stream ciphers," in *Recent Trends in Cryptography*, ser. Contemporary Mathematics, I. Luengo, Ed. Providence, RI, USA: American Mathematical Society, 2009, vol. 477, pp. 21–45.

[68] *3GPP TS 36.211 Release 14*, 3GPP Std., Sep. 2018.

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

[70] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge Univeristy Press, Jul. 2005.

[71] J.-J. van de Beek, M. Sandell, M. Isaksson, and P. O. Börjesson, “Low-complex frame synchronization in OFDM systems,” in *Proc. IEEE Int. Conf. on Univ. Pers. Commun. (ICUPC)*, Tokyo, Japan, Nov. 1995, pp. 982–986.

[72] J.-J. van de Beek, P. O. Börjesson, M.-L. Boucheret, D. Landström, J. M. Arenas, P. Ödling, C. Östberg, M. Wahlqvist, and S. K. Wilson, “A time and frequency synchronization scheme for multiuser OFDM,” *IEEE J. Sel. Areas Commun.*, vol. 17, no. 11, pp. 1900–1914, Nov. 1999.

[73] D. Landström, S. K. Wilson, J.-J. van de Beek, P. Ödling, and P. O. Börjesson, “Symbol time offset estimation in coherent OFDM systems,” *IEEE Trans. Commun.*, vol. 50, no. 4, pp. 545–549, Apr. 2002.

[74] H. Minn, M. Zeng, and V. K. Bhargava, “On timing offset estimation for OFDM systems,” *IEEE Commun. Lett.*, vol. 4, no. 7, pp. 242–244, Jul. 2000.

[75] R. Frank, S. Zadoff, and R. Heimiller, “Phase shift pulse codes with good periodic correlation properties (corresp.),” *IEEE Transactions on Information Theory*, vol. 8, no. 6, pp. 381–382, Oct. 1962.

[76] R. Frank, “Polyphase codes with good nonperiodic correlation properties,” *IEEE Trans. Inf. Theory*, vol. 9, no. 1, pp. 43–45, Jan. 1963.

[77] M. M. U. Gul, S. Lee, and X. Ma, “Robust synchronization for OFDM employing Zadoff-Chu sequence,” in *IEEE Conf. on Inf. Sci. and Sys. (CISS)*. Princeton, NJ, USA: IEEE, Mar. 2012.

[78] M. M. U. Gul, X. Ma, and S. Lee, “Timing and frequency synchronization for OFDM downlink transmissions using Zadoff-Chu sequences,” *IEEE Trans. Wireless Commun.*, vol. 14, no. 3, pp. 1716 –1729, Mar. 2015.

[79] F. Tufvesson, O. Edfors, and M. Faulkner, “Time and frequency synchronization for OFDM using PN-sequence preambles,” in *Proc. IEEE Veh. Tech. Conf. (VTC)*. Amsterdam, The Netherlands: IEEE, Sep. 1999, pp. 2203–2207.

[80] K. Manolakis, D. M. G. Estevez, V. Jungnickel, W. Xu, and C. Drewes, “A closed concept for synchronization and cell search in 3GPP LTE systems,” in *IEEE Wirel. Commun. and Netw. (WCNC)*. Budapest, Hungary: IEEE, Apr. 2009, pp. 1–6.

[81] Q. Wang, C. Mehlfuhrer, C. Mehlführer, and M. Rupp, “Carrier frequency synchronization in the downlink of 3GPP LTE,” in *IEEE Int. Symp. on Pers., Ind. Mob. Rad. Commun. (PIMRC)*. Instanbul, Turkey: IEEE, Sep. 2010, pp. 939–944.

[82] M. Morelli and M. Moretti, “ML estimation of timing, integer frequency and primary sequence index in LTE systems,” in *IEEE Int. Conf. on Commun. (ICC)*. London, UK: IEEE, Jun. 2015, pp. 2609–2613.