Quantized Guessing Random Additive Noise Decoding - A Universal Quantized Soft-Decoder

by

Evan Gabhart

B.S. Electrical Engineering and Computer Science Massachusetts Institute of Technology, 2021

SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF ENGINEERING IN ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

AT THE

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

SEPTEMBER 2022

Author . . . .

DEPARTMENT OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

August 12, 2022

Certified by . . . .

Muriel Médard

Cecil H. Green Professor of Electrical Engineering and Computer Science

Thesis Supervisor

Certified by . . . .

Ken R. Duffy

Director of the Hamilton Institute, Maynooth University

Thesis Supervisor

Accepted by . . . .

Katrina LaCurts

Chair, Master of Engineering Thesis Committee

Quantized Guessing Random Additive Noise Decoding - A Universal Quantized Soft-Decoder

by

Evan Gabhart

Submitted to the Department of Electrical Engineering and Computer Science
on August 12, 2022, in partial fulfillment of the
requirements for the degree of
Master of Engineering of Electrical Engineering and Computer Science

Abstract

Guessing Random Additive Noise Decoding (GRAND) has proven to be a universal, maximum likelihood decoder. Multiple extensions of GRAND have been introduced, giving way to a class of universal decoders. GRAND itself describes a hard-detection decoder, so a natural extension was to incorporate the use of soft-information. The result was Soft Guessing Random Additive Noise Decoding (SGRAND). SGRAND assumes access to complete soft information, proving itself to be a maximum-likelihood soft-detection decoder. Physical limitations, however, prevent one from having access to perfect soft-information in practice.

This thesis proposes an approximation to the optimal performance of SGRAND, Quantized Guessing Random Additive Noise Decoding (QGRAND). I describe the algorithm and evaluate its performance compared to hard-detection GRAND, SGRAND, and another approach to approximating SGRAND, Ordered Reliability Bits GRAND (ORBGRAND). QGRAND also allows itself to be tailored to an arbitrary number of bits of soft information, and I will show as the number of bits increases so does performance. I then use the GRAND algorithms discussed in order to evaluate error correction potential of different channel codes, particularly Polar Adjusted Convolutional codes, CA-Polar codes, and CRCs.

Thesis Supervisor: Muriel Médard
Title: Cecil H. Green Professor of Electrical Engineering and Computer Science

Thesis Supervisor: Ken R. Duffy
Title: Director of the Hamilton Institute, Maynooth University

Acknowledgments

My journey at MIT has been anything but according to plan. At each stage I would have in mind a rigid course, starting with my initial goal in pursuing physics. I am, however, grateful that each twist and turn has led me to where I am now. With each struggle, change of major, and newfound opportunity, I am immeasurably fortunate to have had unwavering support and guidance that has helped me in navigating not only academics but life.

I first would like to sincerely thank my thesis supervisor and research supervisor Professor Muriel Médard for taking me on as an advisee and offering her near limitless wisdom and support. I came to her with little more than eagerness and a desire to learn and in turn have received priceless experiences and a love for learning. Through her guidance I gained not only a wealth of knowledge in the field of communications but also a role model for what being a scholar and engineer should resemble.

I would like to thank Professor Ken R. Duffy who served as my co-advisor for my thesis, offering his vast experience in writing and conducting research. I have been inspired by his ability to use his knowledge in mathematics to approach a wide breadth of problems in new and innovative ways. From being willing to take time with me on a whiteboard to offering life lessons through an interesting anecdote, thank you for your mentorship.

To my academic advisor, Professor Tomás Lozano-Pérez, thank you for your years of advice and support. You always went above and beyond as an advisor and helped me navigate academia and find my way to my current field of study. There were times where I was unsure I would reach the end to this journey but you remained a voice of reason and guidance.

I greatly appreciate the unconditional love and support of my family throughout my education. To my mother, father, and bonus mother, Katie, thank you. It is thanks to you all that I have been instilled with perseverance, a strong sense of honor, and an appreciation for learning. These traits have proven invaluable as I navigate life and without them my accomplishments would not be possible.

To my partner, Sahara, thank you for your unwavering love and encouragement. In our time together you have always been there to provide comfort during struggles and celebration during victories. You have taught me to advocate for myself, the importance of having a good laugh, and engineering should be used to make something fun. I am grateful to have you at my side.

The project depicted was or is sponsored by the Defense Advanced Research Projects Agency under grant number HR00112120008, the content of the information does not necessarily reflect the position or the policy of the government, and no official endorsement should be inferred.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

At any point that digital data is transmitted, one runs the risk of errors between when the data is sent and is received. In order to mitigate this risk, one introduces redundancy in the form of coding. Encoding a message can be thought of as pooling risk across the bits sent in hopes enough information is retained to recover the original message. Traditionally this involves a designated encoder and decoder pair, with the decoder being designed to take advantage of the structure of the code. This strict coupling of encoder and decoder restricts the forms of detection available and the hardware used. Guessing Random Additive Noise Decoding (GRAND)[1], however, has led to a class of universal decoders that offer a Maximum Likelihood decoding by taking a noise-centric approach. GRAND decoders guess noise effects in decreasing order of likelihood, subtracting this noise from the received signal, and checking for code-book membership. The first valid code-word then is an ML-decoding.

Soft-Guessing Random Additive Noise Decoding (SGRAND)[2] applies this approach with access to full soft information, providing an optimal benchmark for error-correction performance. In practice, however, real-valued soft information is less feasible. Additionally SGRAND makes use of dynamic memory in the form of a max heap, the complexity of which is not suitable for hardware. The problem then lies in approximating the performance of SGRAND given a limited amount of soft information in a way that is circuit-friendly.

This thesis presents Quantized Guessing Random Additive Noise Decoding, or

QGRAND. QGRAND quantizes soft information into one of a number of discrete values for an arbitrary number of bits of soft information available. This allows QGRAND to be tailored to application need and provides performance that approaches that of SGRAND as the number of quantization bins increases. I evaluate QGRAND by simulating additive white Gaussian noise (AWGN) channels and observing bit error rate (BER) and block error rate (BLER). These results are compared against those of hard-detection GRAND, SGRAND, and the current state of the art approximation to SGRAND, Ordered Reliability Bits GRAND (ORBGRAND)[3].

1.1 Thesis Outline

This thesis begins with Chapter 2 covering relevant background and related work regarding decoding in communications and the noise-centric approach offered by GRAND algorithms. Chapter 3 will formally describe QGRAND algorithm and provide a clear definition as well as figures to assist in conceptualization. Chapter 4 will evaluate performance with regards to error correction. This will be for a variety of codes of different rates compared to other GRAND variants. Chapter 5 evaluates different codes inherent potential for error correction when decoded using GRAND algorithms in the high-rate, short block-length setting. Chapter 6 gives conclusions and closing remarks.

Chapter 2

Background

2.1 GRAND

The increasing demand for ultra reliable low latency communications (URLLC)[4] has encouraged research of decoders with high accuracy at short block lengths. GRAND decoders are universal, maximum likelihood decoders that rather than take advantage of the structure of a channel code, they take an agnostic approach that only considers the received signal and its proximity to a valid code-word. Whereas a traditional maximum likelihood decoder would determine the likelihood of possible code-words $c \in C$ and guess them from most likely to least likely, GRAND considers the equivalent problem of the noise effect. For a demodulated channel output $y = demod(Y)$, we define a received signal $z = c \oplus y$ for sent code-word c. Then we have $c = c = z \oplus y$, so determining c amounts to guessing the value of z. GRAND was first proposed for hard-detection over an AWGN channel where this problem amounts to guessing in order of increasing hamming weight, starting with the zero vector. Other variants such as Soft GRAND make use of soft information in order to improve performance. The GRAND Algorithm is described in Figure 2-1. GRAND in its hard-detection form already has exhibited its ability to be implemented in circuits[5][6].

Input: Code-book membership function $C : \{0,1\}^n \mapsto \{0,1\}$; demodulated bits y^n; Optional Information Φ.
Output: Decoding $c^{*,n}$.

```
d ← 0
while d=0 do
    z^n ← next most likely noise effect vector which may depend on Φ
    D ← 1
    if C(y^n ⊕ z^n) = 1 then
        c^{*,n} ← y^n ⊕ z^n
        d ← 1
        return c^{*,n}
    end if
end while
```

Figure 2-1: Guessing Random Additive Noise Decoding algorithm

2.2 Soft-Detection and GRAND

SGRAND is a GRAND decoder that avails itself of full soft information. SGRAND then can organize noise effects by absolute likelihood within a max-heap data structure. SGRAND has proven itself to be a maximum likelihood soft decoder, meaning it provides optimal performance in the soft-detection setting. SGRAND accomplishes this by utilizing the real valued bit reliabilities of a received signal Y. Let reliability be defined as

$$L(Y) = LLR(|Y|) = \left|\log\left(\frac{f_{Y|C}(Y|1)}{f_{Y|C}(Y|0)}\right)\right| \tag{2.1}$$

Then SGRAND determines noise effects in increasing order of $\sum_{i=1}^{n} L(Y_i)z_i$, equivalent to guessing noise effects in decreasing order of likelihood. The complexity, however, does not lend itself well to implementation in circuits. This has inspired GRAND algorithms that are able to approximate the performance of SGRAND given limited soft information. A goal of a number of GRAND algorithms then is to approximate $L(Y)$. Let these approximations be denoted by $\widehat{L(Y)}$. The current state-of-the-art approach to this is Ordered Reliability Bit GRAND. ORBGRAND sorts bits by increasing order of reliability. A logistic weight, defined as the index of a particular bit under this reliability ordering, is assigned to each bit. Consider the case of hard-detection GRAND where noise effects are sorted by increasing hamming weight. ORBGRAND instead uses this logistic weight as a substitute for hamming weight, ordering noise effect guesses by increasing order of total summed logistic weight. Letting this logistic weight serve as $\widehat{L(Y)}$ ORBGRAND then queries noise effects by increasing order of $\sum_{i=1}^{n} \widehat{L(Y_i)} z_i$. For example, consider a three bit sequence already in order of increasing reliability. Then the first bit has logistic weight 1, the second has logistic weight 2, and the third logistic weight 3. The first five noise effects to be guessed then would be:

Query No.	Noise Effect
1	(0, 0, 0)
2	(1, 0, 0)
3	(0, 1, 0)
4	(1, 1, 0)
5	(0, 0, 1)

Table 2.1: Example of ORBGRAND Query Order

with ties broken arbitrarily. One can observe in order summed logistic weights of 0, 1, 2, 3, and 3 for these noise effects.

Chapter 3

QGRAND

ORBGRAND relies on $\log(n)$ bits of soft information in order to sort received bits by order of reliability. The sorting operation itself also comes at some computational cost. In addition, performance of ORBGRAND at high signal-to-noise ratio produces a less than optimal guessing order as the difference in reliability between bits becomes very small. This inspired a look at an alternative quantization of soft information in the form of binning together bits of the same order of reliability. Then as the difference in reliability of bits becomes small, the guessing order is likely to reflect this in more bits being binned together as roughly equally likely to be erroneous. The culmination of this became Quantized Guessing Random Additive Noise Decoding (QGRAND).

3.1 QGRAND Algorithm

QGRAND is able to be tailored to an arbitrary number of bits of available soft information. Let the number of bits of soft information available be q. Then QGRAND is able to assign received bits to $Q = 2^q$ distinct bins dependent on real valued reliabilities. Let each reliability value be associated with a bin with distinct weight via function $W : Y_i \mapsto \mathbb{N}$. From this we have $\widehat{L(Y)} = \sum_{i=1}^{n} W(Y_i) * z_i$. The algorithm is then described in Figure 3-1.

Input: Code-book membership function $C : \{0,1\}^n \mapsto \{0,1\}$; function mapping reliability values to weights $W : \mathbb{R} \mapsto \mathbb{N}$; real valued reliability vector $L(Y^n)$; demodulated signal y^n; number of quantization bins Q.
Output: Decoding $c^{*,n}$

```
d ← 0, w ← 0
while d = 0 do
    for z^n s.t. Σ_{i=1}^n W(L(Y_i)) z_i = w do
        if C(y^n ⊕ z^n) = 1 then
            c^{*,n} ← y^n ⊕ z^n
            d ← 1
            return c^{*,n}
        end if
    end for
    w ← w + 1
end while
```

Figure 3-1: QGRAND Algorithm. Takes code-book membership function, weight assignment function, real valued reliability of bits, demodulated signal, and number of quantization bins as inputs. Outputs decoding.

3.2 Exploring a Quantization

The first attempt at a quantization considered bins determined by orders of magnitude. One assigns bits with a likelihood of error on the order of 10^{-1} to bin 1, the order of 10^{-2} to bin 2, etc. Then bin 1 is associated with a weight of 1, bin 2 a weight of 2, and so on. Then noise effects are to be guessed in increasing order of summed weight. For example, assume we have 2 bins, one for order of 10^{-1} and another bin for less erroneous bits. Consider a 3 bits message of weights 1, 2, and 1 for the first, second, and third bits respectively. The first five queried noise effects are as follows:

Query No.	Noise Effect
1	(0, 0, 0)
2	(1, 0, 0)
3	(0, 0, 1)
4	(1, 0, 1)
5	(0, 1, 0)

Table 3.1: Example of QGRAND Query Order

with ties broken arbitrarily.

This preliminary quantization, however, proved sub-optimal and did not scale well for varying SNR and number of quantization bins. A bin size dependent on the standard of deviation of the noise σ and the number of quantization bins Q then was constructed. Consider quantizing the real valued reliabilities $L(Y) = LLR(|Y|)$ to natural numbers 1 to Q. The bin size was defined as $\frac{1-\sigma/2}{Q}$. Reliability values ranging from 0 to the bin size are associated with a bin of weight 1, values between the bin size and twice the bin size are associated with a bin of weight 2, etc. This expression of bin size was determined empirically under the assumption of using natural numbers from 1 to Q as weights and constant bin size.

However, if one increases the width of more reliable bins relative to those of low reliability, this effectively increases resolution of our quantization for the most likely to be flipped bits. Let bin sizes be determined by a function of a value β. The previously mentioned binning used $\beta = \frac{1-\sigma/2}{Q}$ with bin ranges $[0, \beta), [\beta, 2\beta), ...[Q\beta - \beta, +\infty)$. An alternative expression for β is proposed after analyzing the achievable rate of different choices of β. We then instead use $\beta = \frac{1-\sigma/2}{2Q-1}$ and a binning determined by ranges $[0, \beta), [\beta, 3\beta), ...[2Q\beta - \beta, +\infty)$. In this new binning, the lowest reliability bin then contains approximately 30% of bits are accurately binned while the remaining 70% are binned together. For results contained herein, the latter choice of β and binning are used.

Chapter 4

QGRAND Performance Evaluation

As the primary objective of QGRAND is to approximate the universal, ML soft decoder SGRAND, SGRAND is used as a performance reference. Additionally as ORBGRAND is the current state-of-the-art GRAND algorithm for quantized soft decoding, ORBGRAND is also used as a benchmark. Also included is hard-detection GRAND. In the following simulations a randomly generated message is encoded according to a chosen channel code and then modulated with binary phase shift keying (BPSK). The modulated signal is then sent over a simulated AWGN with noise of amplitude 1 and variance $\sigma^2 = \frac{1}{10^{\mathrm{SNR}/10}}$ where SNR is measured in dB and is calculated by $\mathrm{SNR} = E_b/N_0 + 10\log_{10}(k/n)$. The noise is randomly generated from this Gaussian distribution and the resulting signal is demodulated. The pre-demodulated signal is also used as soft-information. In the case of QGRAND, a number of quantization bits (q) are specified, where 2^q is the number of quantization bins Q. If the decoding does not match the original code-word then the transmission is recorded as an error and the number of erroneous bits are also recorded. This is repeated to determine performance statistics such as block error rate and bit error rate by averaging over all iterations.

4.1 BLER

These BLER results compare QGRAND with different numbers of quantization bits versus that of hard-detection GRAND, the state-of-the-art quantized soft GRAND decoder ORBGRAND, and the ML soft decoder SGRAND. Figure 4-1 shows BLER performance for CA-Polar[7] codes of rate R $= n/k = 0.91$. Also included is the performance of CA-SCL decoding with list size of 32[8][9], as this is the standard method of decoding for CA-Polar codes. The gain for each additional bit of quantization ranges from 0.25 - 0.5 dB. Additionally at 6 bits of soft information QGRAND shows marginal BLER performance over that of ORBGRAND, with ORBGRAND using $\log_2(n) = 7$ bits of soft information for its rank-ordering.

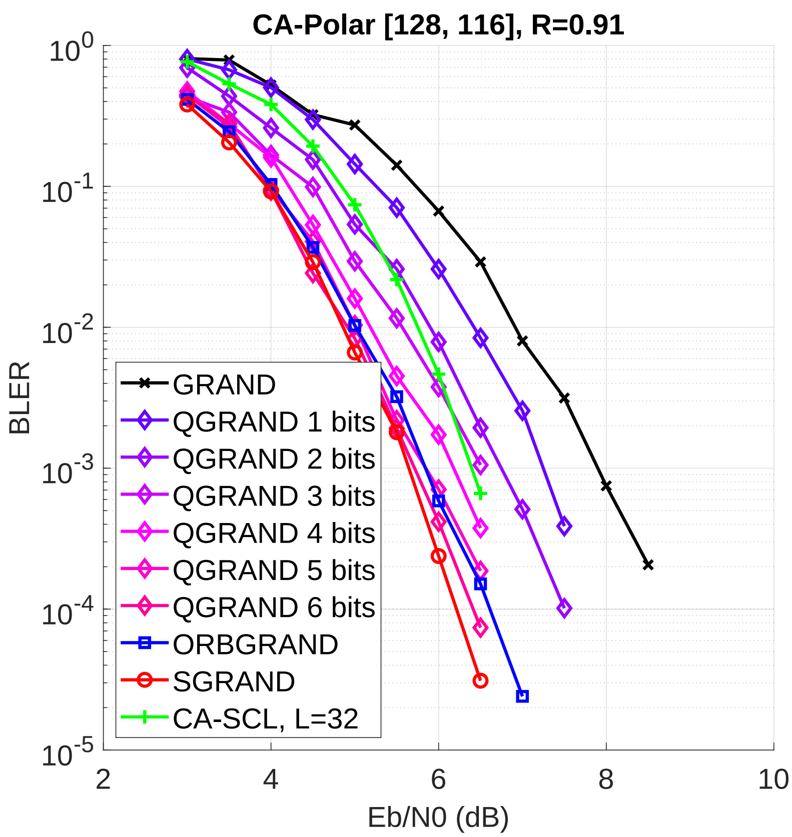

Figure 4-1: BLER performance of QGRAND and other GRAND decoders for CA-Polar [128,116] code

Figure 4-2 shows results for a cyclic redundancy check (CRC) code of rate R=0.95 as GRAND algorithms can use them for error correction as opposed to just detection. Similar to previous, again shows marginal gains over ORBGRAND at 6 bits of soft information and improvements in performance as the amount of soft information increases.

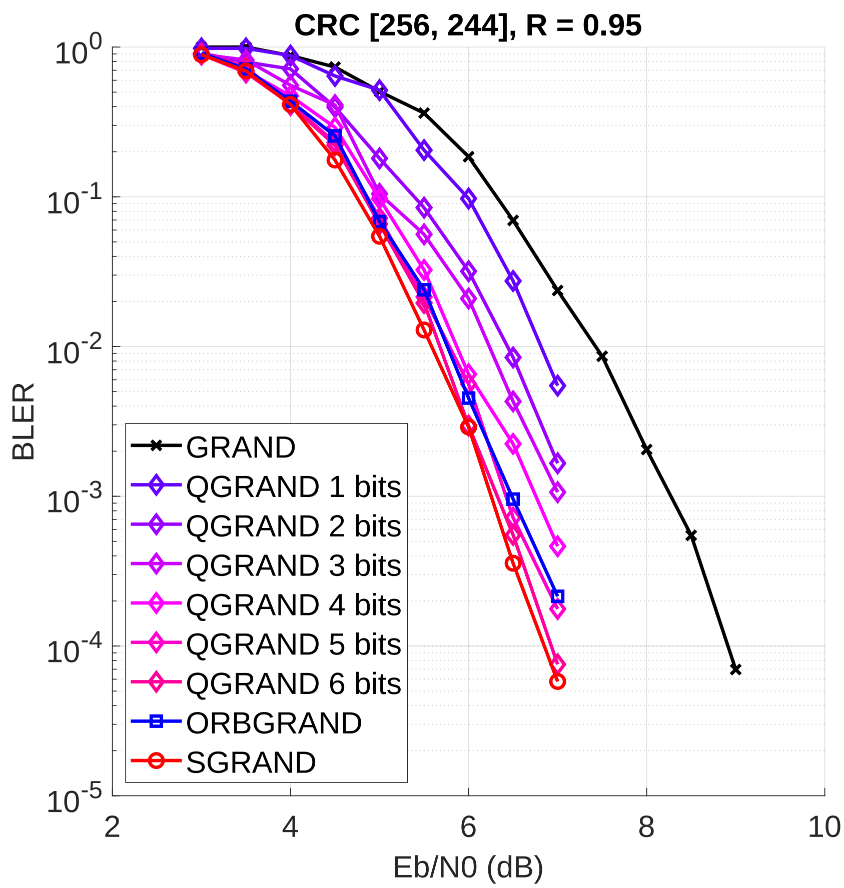

Figure 4-2: BLER performance of QGRAND and other GRAND decoders for CRC [256,244] code

4.2 BER

Bit error rate results demonstrate the ratio of erroneous bits received per transmission. Bit errors are determined by comparing the original k bit message to the first k

decoded bits, counting the number of differing bits. Figure 4-3 shows BER results for a CA-Polar code of rate R = 0.95. The results show again better performance as soft information increases, this time with both 5 and 6 bits providing performance competitive to that of ORBGRAND.

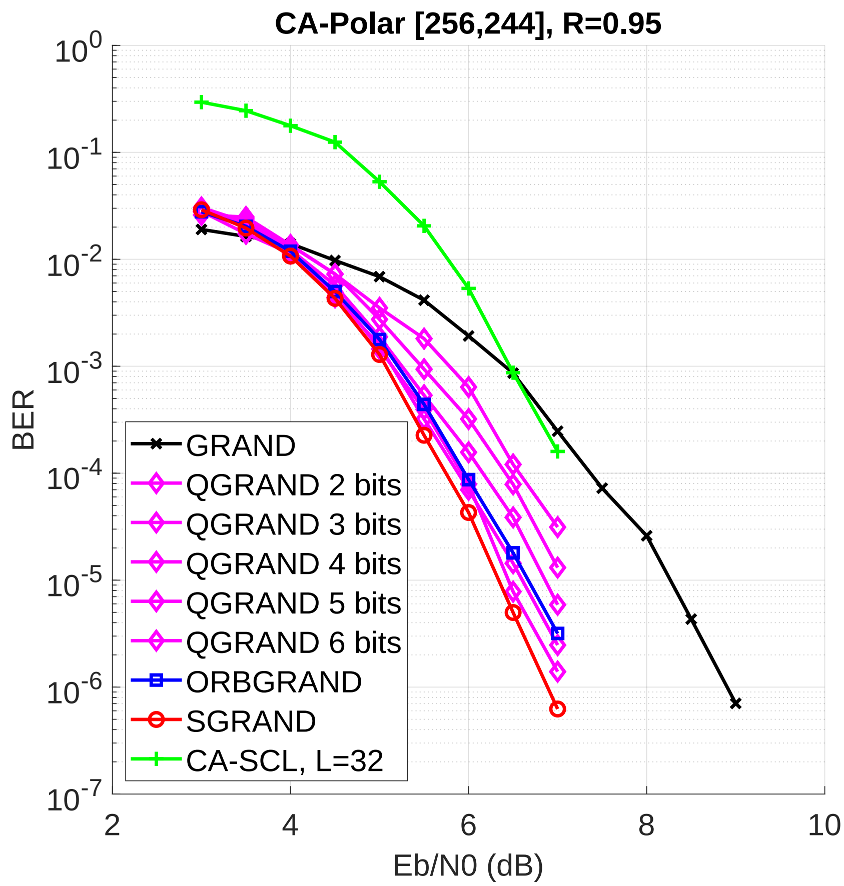

Figure 4-3: BER performance of QGRAND and other GRAND decoders for CA-Polar [256,244] code

Figure 4-4 depicts BER results for Bose-Chaudhuri-Hocquenghem (BCH) codes of rate 0.94. In this case of a higher rate, short block length code 3 quantization bits show to be competitive to the bit error rate of ORBGRAND at the noise levels shown.

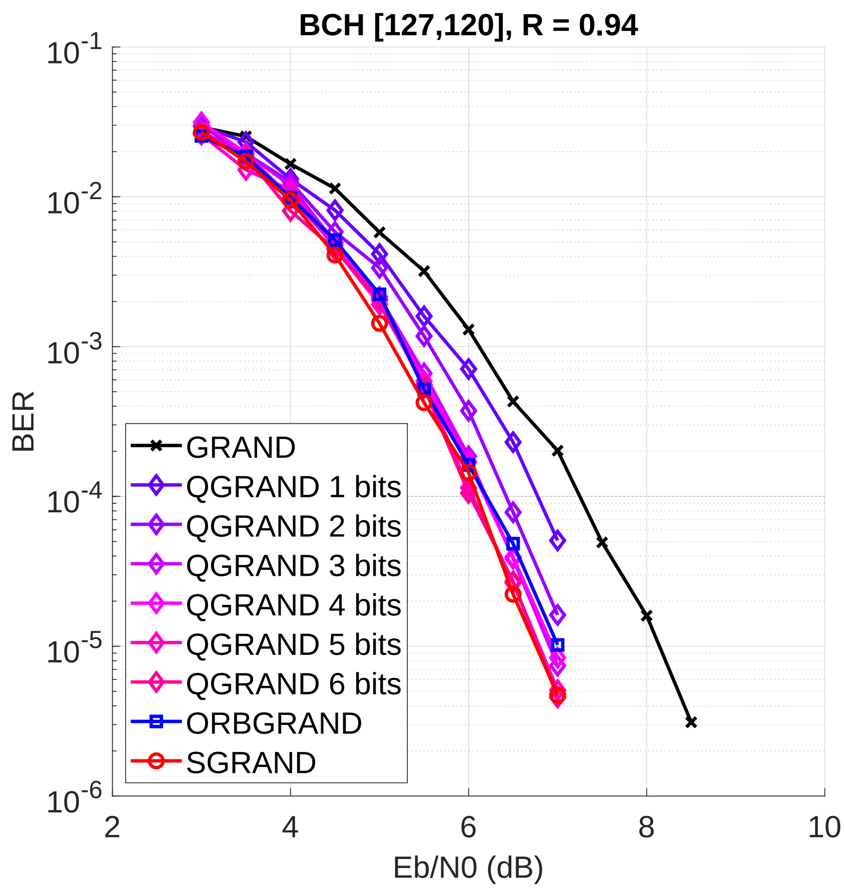

Figure 4-4: BER performance of QGRAND and other GRAND decoders for BCH [127,120] code

Chapter 5

Inherent Error-Correction Capability of Codes

A number of emerging technologies have increased demand for ultra reliable, low-latency communications. Traditionally, high reliability codes involve structuring the code and its decoder as a pair to take advantage of the code's structure for error-correction. However URLLC has increased interest in codes at shorter block lengths in order to reduce latency, renewing interest in finding the most suitable codes. GRAND has shown itself to be a universal, maximum likelihood decoder. Since GRAND can be used for any block code, this puts in a unique position to evaluate the potential of different codes for their capacity for error-correction. This gives an unbiased assessment as the code-agnostic nature of GRAND separates the error correction abilities of codes from their usual decoders and being an ML decoder ensures the performance is optimal.

5.1 PAC codes

Polar codes have shown to be the first explicit family of capacity achieving codes for a binary symmetric channel. However, polar codes are known on their own to have performance suffer at short block lengths. In order to mitigate this, a state of the art solution is to pair polar codes with a cyclic redundancy check prior to polar coding, with these being known as CRC Assisted Polar (CA-Polar) codes. The improvements to performance resulted CA-Polar codes being the standard for 5G-NR communications. Recently, as opposed to a CRC pre-coding, a convolutional pre-coding has been proposed. This is known as polarization-adjusted convolutional (PAC) codes[10]. PAC codes have been shown with the use of a Fano decoder[11], a type of sequential decoder, to offer improvements to error-correction performance over CA-Polar codes.

5.2 Performance Evaluation using ML Decoder

In the performance evaluation of PAC codes in [10] a Fano decoder is used while CA-Polar codes made use of a CA-SCL decoder. A decoder-independent approach to comparing their performance would be to use the same universal decoder for each. Otherwise this observed difference in performance could be partially due to how well each decoder takes advantage of the code structure. Since the product of two linear block codes is itself a linear block code, this allows us to use GRAND in order to evaluate the performance. SGRAND being a maximum-likelihood decoder availing of full soft-information makes it an optimal performance benchmark for PAC codes. The error-correction performance is compared to that of SGRAND used for a CA-Polar code and a CRC all of rate R $= k/n = 0.95$. Figure 5-1 shows that for each of these the bit error rate is identical. When one divorces these codes from their paired decoders, one is left by the error correction potential offered by the code itself. The CRC code performing just as well is particularly of interest as a CRC is traditionally only used for error detection. A CRC is also able to be of any length as opposed to polar coding requiring a power of 2 length.

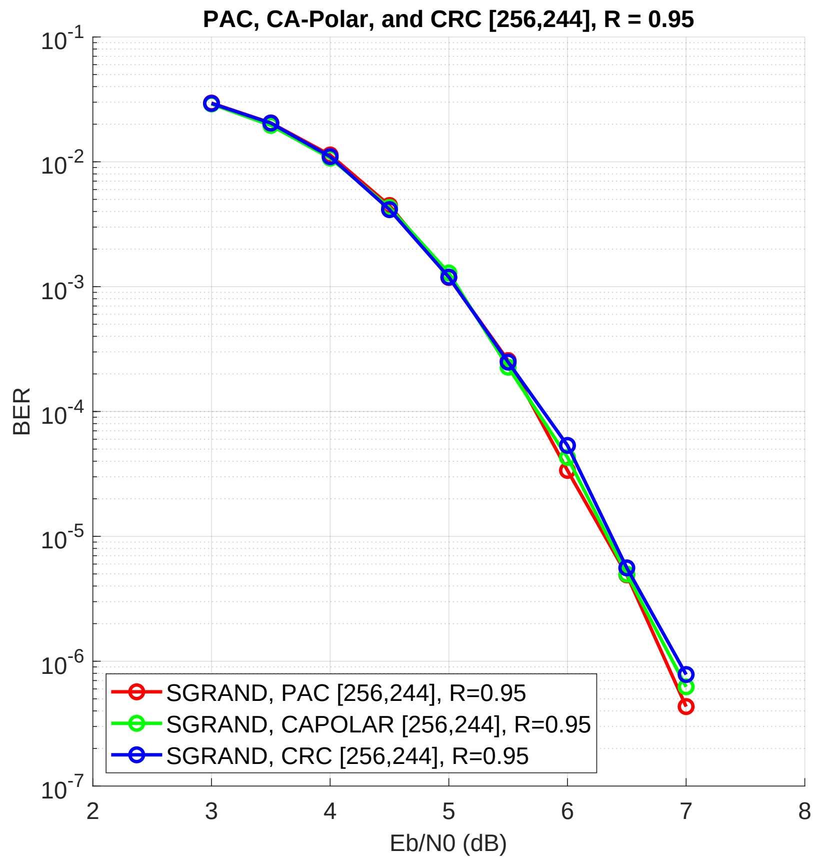

Figure 5-1: BER performance of PAC, CA-Polar, and CRC codes of rate 0.95 with SGRAND

Figure 5-2 depicts the results for block error rate of these same codes. This performance reinforces what is observed in the BER results. The same simulation is also applied to rate $R = 0.95$ codes with synonymous results. For high rate codes of short block length, GRAND algorithms show themselves to be a useful tool in evaluating the potential for error correction. In particular they are useful for observing performance separate from a code's paired decoder.

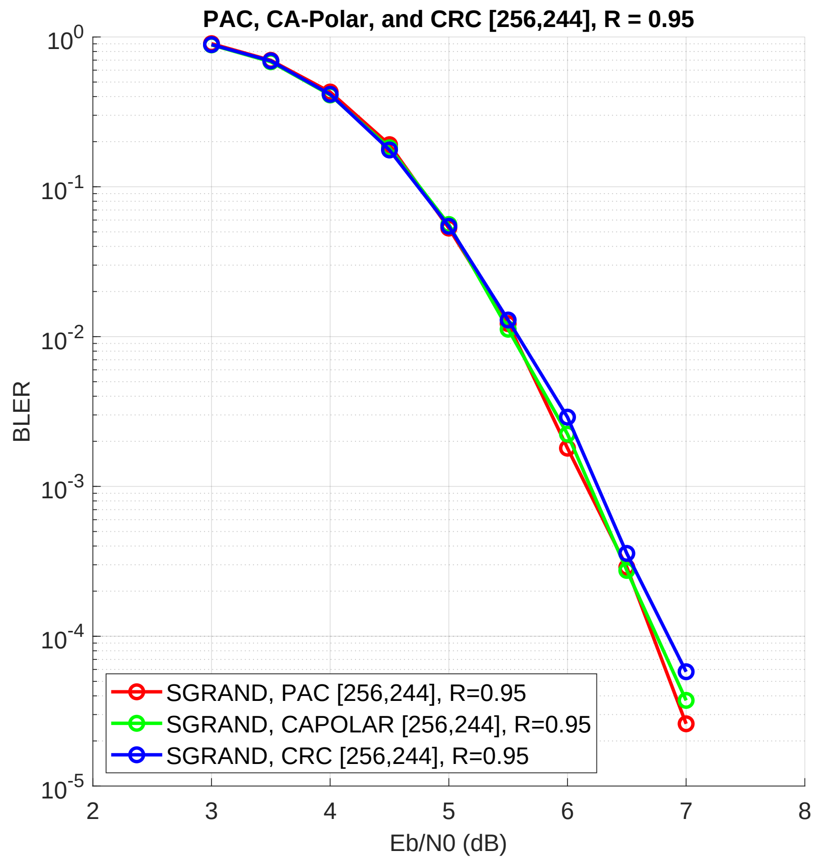

Figure 5-2: BLER performance of PAC, CA-Polar, and CRC codes of rate 0.95 with SGRAND

5.3 Performance Evaluation using Quantized Soft Decoders

While SGRAND has shown these codes to perform just as well as one another for high rate and short block lengths, it can be useful to look at quantized soft decoders. Perhaps a particular code may make better use of incomplete soft information, or has advantages at a particular amount or type of soft information. As our tools for investigating this scenario we turn to QGRAND and ORBGRAND. QGRAND's ability to be tailored to an arbitrary number of quantization bits allows one to observe the error correction performance across these codes as the number of quantization bits changes. Figures 5-3 and 5-4 show the bit error rate performance of each of PAC, CA-Polar, and CRC codes for 1 through 6 bits of soft information at rates R = 0.91 and R = 0.95 respectively. It is observed that for a given number of quantization bits, QGRAND offers identical performance across each of these codes for a fixed rate. This shows that the improvements to performance offered by increased soft information with QGRAND is agnostic to each of these codes.

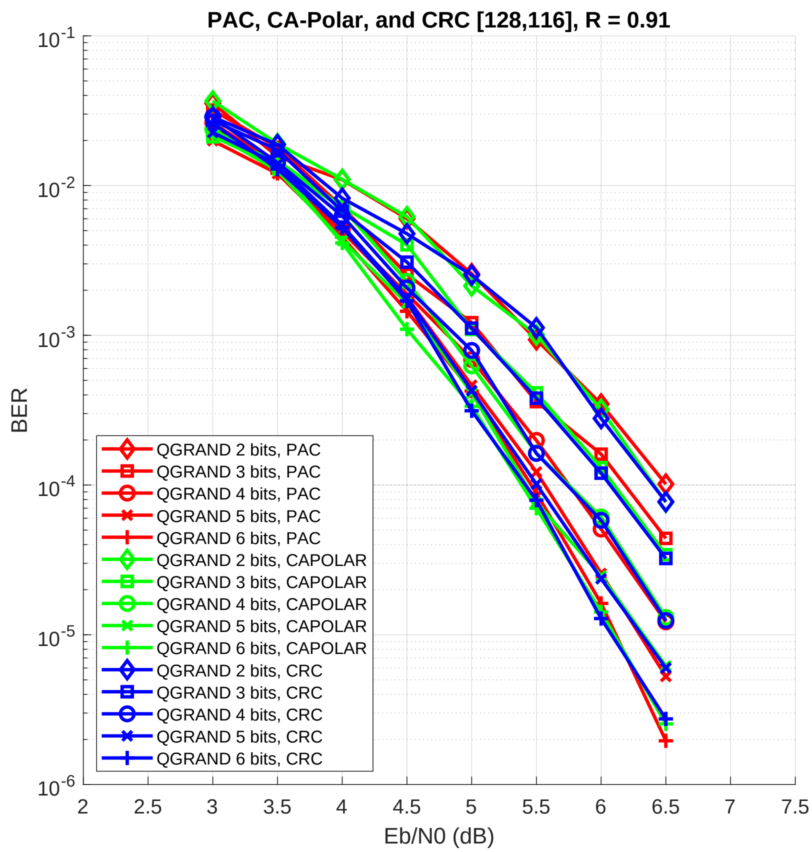

Figure 5-3: BER performance of PAC, CA-Polar, and CRC codes of rate 0.91 with QGRAND

Figure 5-5 compares the bit error rate performance of each of these codes of the same rate using ORBGRAND for rate R = 0.95, again showing identical performance.

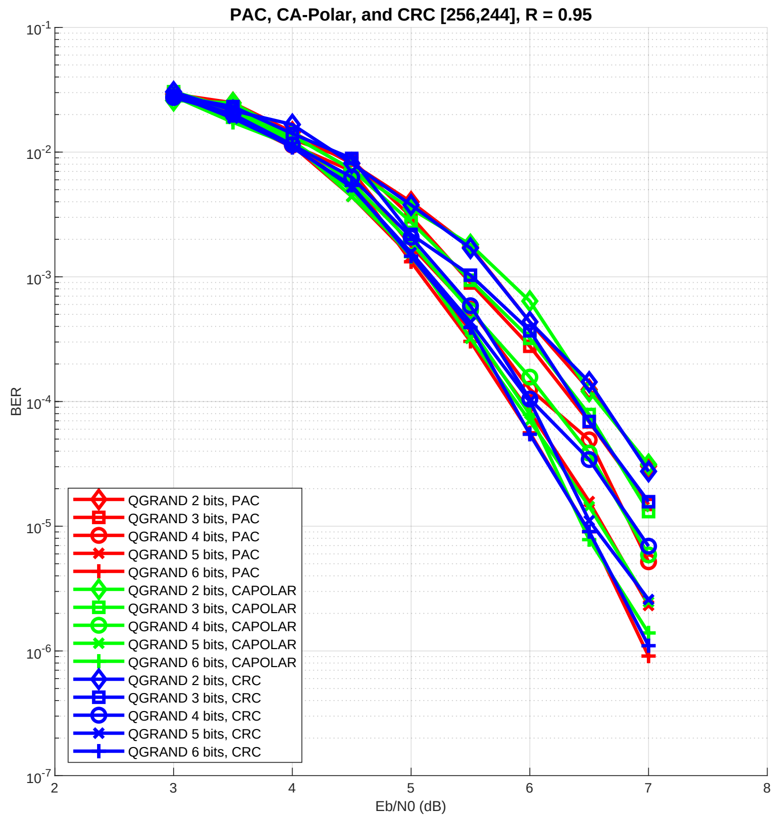

Figure 5-4: BER performance of PAC, CA-Polar, and CRC codes of rate 0.95 with QGRAND

Synonymous results are also achieved for block error rate. This shows the same inherent ability for these codes to be used by ORBGRAND and QGRAND for error correction, at least for the high-rate, short block-length scenario.

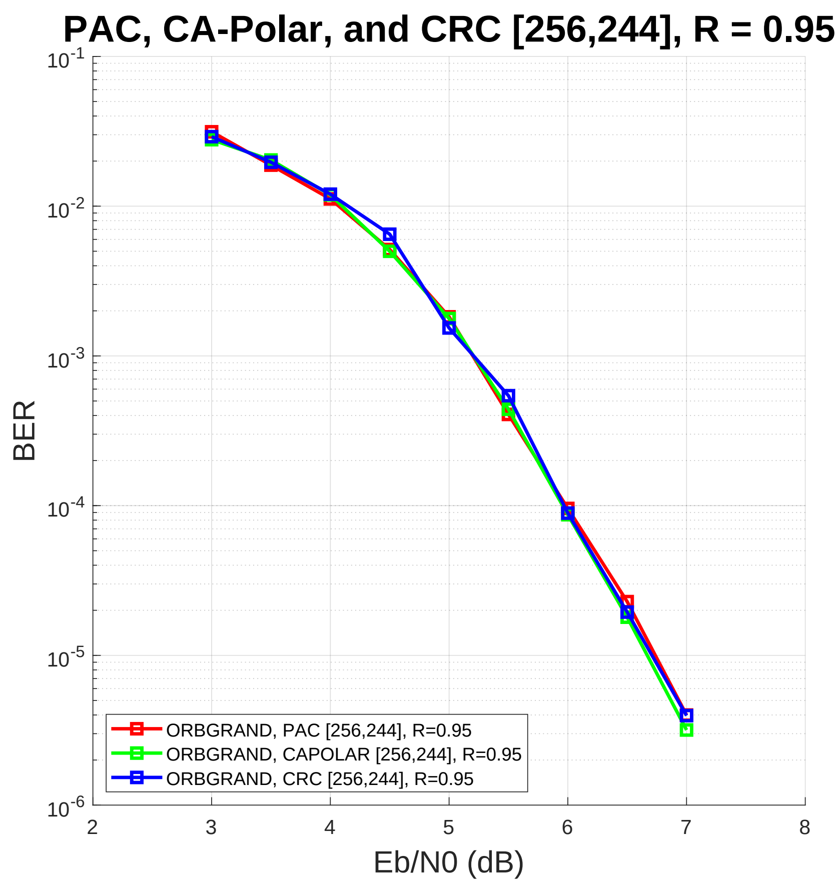

Figure 5-5: BER performance of PAC, CA-Polar, and CRC codes of rate 0.95 with ORBGRAND

Therefore, although PAC codes with a sequential Fano decoder offers performance advantages over CA-Polar codes using CA-SCL, this can likely be attributed a great deal to the decoder used.

Chapter 6

Conclusions

GRAND algorithms offer a unique approach to decoding that has provided a class of universal decoders. With the demand for increased reliability, soft-detection is essential. SGRAND has shown itself to be a code agnostic, ML soft decoder. For practicality, it is of interest to determine how best to quantize soft information in order to approximate the ideal performance offered by SGRAND. This thesis shows QGRAND to provide error correction performance rivaling that of the state-of-the-art GRAND decoder for soft-detection, ORBGRAND. With the universality inherent to GRAND decoders and the ability to be tailored to an arbitrary amount of soft information, QGRAND exhibits flexibility to be applied to a wide range of applications.

Additionally, GRAND algorithms in general prove a useful tool for evaluating the error-correction potential of different channel codes. As traditionally a channel code is paired with a decoder that is tailored to the code's structure, it is difficult to give a fair assessment of relative error-correction performance. The code-book independent nature of GRAND allows one to divorce this partnership of code and decoder, letting performance be determined only by the qualities of the code. Additionally, with SGRAND proving itself an ML soft decoder, SGRAND allows one to observe the optimal error correction performance of a given code. Through the use of these tools this thesis has shown in high-rate, short block length settings that CRC codes can provide error-correction just as effectively as PAC codes and CA-Polar codes when used with GRAND, with the added benefit of not being restricted to power of 2 block lengths. As more technologies increase demand for URLLC, GRAND decoders can prove useful in determining the effectiveness of particular codes and give an unbiased evaluation of their potential.

6.1 Future Work

Future work regarding QGRAND would involve a more circuit-friendly, parallelized implementation. For the purposes of evaluation and simulation a sequential version of QGRAND is used that is currently bottle-necked by a combinatorics problem of integer partitions of multi-sets with limited repetition. This is due to the generation of noise effects of a given weight equating to determining all valid partitions of that weight given a multi-set of the number of bits belonging to each bin. Perhaps this could be mitigated via a dynamic programming approach to sequentially generate valid combinations for a given weight, or an algorithm for traversing adjacent branches of a tree of noise effects in parallel.

Bibliography

[1] K. R Duffy, J. Li, and M. Médard. Guessing noise, not code-words. In *IEEE Int. Symp. on Inf. Theory*, 2018.

[2] A. Solomon, K. R. Duffy, and M. Médard. Soft maximum likelihood decoding using GRAND. In *IEEE ICC*, 2020.

[3] K. R. Duffy. Ordered reliability bits guessing random additive noise decoding. In *IEEE ICASSP*, pages 8268–8272, 2021.

[4] Giuseppe Durisi, Tobias Koch, and Petar Popovski. Toward massive, ultrareliable, and low-latency wireless communication with short packets. *Proc. IEEE*, 104(9):1711–1726, 2016.

[5] A. Riaz, V. Bansal, A. Solomon, W. An, Q. Liu, K. Galligan, K. R. Duffy, M. Médard, and R. T. Yazicigil. Multi-code multi-rate universal maximum likelihood decoder using GRAND. In *IEEE ESSCIRC*, 2021.

[6] Arslan Riaz, Muriel Medard, Ken R Duffy, and Rabia Tugce Yazicigil. A universal maximum likelihood grand decoder in 40nm CMOS. In *COMSNETS*, pages 421–423, 2022.

[7] K. Niu and K. Chen. CRC-aided decoding of Polar codes. *IEEE Commun. Lett.*, 16(10):1668–1671, October 2012.

[8] I. Tal and A. Vardy. List decoding of Polar codes. In *IEEE Int. Symp. Inf. Theory*, pages 1–5, 2011.

[9] A. Balatsoukas-Stimming, M. B. Parizi, and A. Burg. LLR-based successive cancellation list decoding of Polar codes. *IEEE Trans. Signal Process.*, 63(19):5165–5179, Oct 2015.

[10] E. Arikan. From sequential decoding to channel polarization and back again. *ArXiv*, abs/1908.09594, 2019.

[11] R. Fano. A heuristic discussion of probabilistic decoding. *IEEE Transactions on Information Theory*, 9(2):64–74, 1963.