

Approaching Maximum Likelihood Decoding Performance via Reshuffling ORBGRAND

Li Wan and Wenyi Zhang
University of Science and Technology of China
Email: wenyizha@ustc.edu.cn

Abstract—Guessing random additive noise decoding (GRAND) is a recently proposed decoding paradigm particularly suitable for codes with short length and high rate. Among its variants, ordered reliability bits GRAND (ORBGRAND) exploits soft information in a simple and effective fashion to schedule its queries, thereby allowing efficient hardware implementation. Compared with maximum likelihood (ML) decoding, however, ORBGRAND still exhibits noticeable performance loss in terms of block error rate (BLER). In order to improve the performance of ORBGRAND while still retaining its amenability to hardware implementation, a new variant of ORBGRAND termed RS-ORBGRAND is proposed, whose basic idea is to reshuffle the queries of ORBGRAND so that the expected number of queries is minimized. Numerical simulations show that RS-ORBGRAND leads to noticeable gains compared with ORBGRAND and its existing variants, and is only 0.1dB away from ML decoding, for BLER as low as 10^{-6} .

Index Terms—Error pattern, GRAND, maximum likelihood decoding, ORBGRAND, querying, search problem

I. INTRODUCTION

In recent years, a variety of decoding methods generally referred to as guessing random additive noise decoding (GRAND) have been proposed [1] [2] [3] [4]. These decoding methods are universal in the sense that they are applicable to all block codes, building upon the key idea of conducting a sequence of queries to test whether a queried error pattern can convert a noise-corrupted channel output vector into a codeword. They are particularly suitable for codes with relatively short length and high rate [1] [5], meeting the requirements of ultra-reliable low-latency communication (URLLC). Among the variants of GRAND, soft GRAND (SGRAND) [4], which schedules queries based on the exact values of log likelihood ratios (LLRs), is in fact equivalent to maximum likelihood (ML) decoding [6] [7, Ch. 10] (assuming that the queries exhaust all possible error patterns, in the worst case). It, however, is not amenable to hardware implementation because its queries are generated based on the exact values of LLRs in a sequential fashion. On the other hand, ordered reliability bits GRAND (ORBGRAND) [3] only needs the relationship of ranking among the magnitudes of LLRs to generate queries. Once the ranking is obtained, the queries can then be implemented in a parallel fashion, thereby facilitating efficient hardware implementation [3] [8] [9] [10] [11].

From an information-theoretic perspective, ORBGRAND has been shown to be almost capacity-achieving [6]. For finite code lengths, however, ORBGRAND still exhibits a noticeable performance loss compared with ML decoding in terms of block error rate (BLER), especially as signal-to-noise ratio (SNR) increases. Several works have focused on improving ORBGRAND while still retaining its amenability to hardware implementation (i.e., using the ranking among magnitudes of LLRs only), by scheduling queries based on various heuristics [3] [6] [11], typically related to the cumulative distribution function (CDF) of the magnitude of LLR. Other enhancement techniques include adopting list decoding [12] and combining ad hoc queries based on empirical evidences [13].

In this work, we propose an improvement of ORBGRAND, termed RS-ORBGRAND, motivated by the insight that, if, ideally, one could reshuffle the queries of ORBGRAND so that they are identical to those of SGRAND, then the performance of the thus reshuffled ORBGRAND would be exactly that of SGRAND and hence be optimal. Such an idealistic situation, of course, is impossible, but it provides a guideline for us, and drives us to study the difference between an arbitrary querying schedule and the querying schedule of SGRAND.

In order to conduct a tractable analysis, we turn to a search problem, wherein a randomly selected codeword is corrupted by additive white Gaussian noise (AWGN) and is observed by a searcher, and the searcher conducts a sequence of queries to recover this particular codeword. Such a problem is an idealized decoding problem, without needing to consider the risk of mistakenly deciding an incorrect codeword. For this search problem, we analyze the expected number of queries, showing that the querying schedule that minimizes the expected number of queries obeys a monotonicity property. Consequently, we propose to reshuffle the querying schedule generated by ORBGRAND so as to satisfy this monotonicity property. Since the reshuffling is offline, this method follows the same decoding process as ORBGRAND and has low complexity. When applying the thus obtained RS-ORBGRAND to decoding, numerical simulations using BCH and polar codes show that it leads to a gain of no less than 0.3dB compared with ORBGRAND and its existing variants, and is only 0.1dB away from ML decoding, for BLER as low as 10^{-6} .

The remaining part of this paper is organized as follows. Section II introduces the channel model and GRAND. Section III analyzes the search problem motivated by GRAND, and proposes RS-ORBGRAND. Section IV applies RS-

ORBGRAND to decoding, and verifies its effectiveness via numerical simulations. Section V concludes this paper. We use uppercase letters (e.g., X) to represent random variables and their corresponding lowercase letters (e.g., x) to represent their realizations.

II. PRELIMINARIES

A. Channel Model

We consider general binary block coding wherein information bits $\underline{U} \in \mathbb{F}_2^k$ generate a codeword $\underline{W} \in \mathcal{C} \subseteq \mathbb{F}_2^n$. The code rate is thus $R = \frac{k}{n}$. For simplicity, in this work we consider BPSK modulation over an AWGN channel, while the basic idea should be applicable, with appropriate generalizations, to general memoryless channels. We assume that each element of \underline{W} , W_i , $i = 1, \dots, n$, takes 0 and 1 with equal probability, a condition commonly satisfied for linear block codes. With BPSK, the transmitted vector \underline{X} is such that $X_i = 1$ if $W_i = 0$ and $X_i = -1$ if $W_i = 1$, $i = 1, \dots, n$; and the resulting AWGN channel output vector satisfies $\underline{Y}|\underline{X} \sim \mathcal{N}(\underline{X}, \sigma^2 \mathbf{I}_{n \times n})$. Given \underline{Y} , the LLRs are

$$L_i = \log \frac{P(Y_i | W_i = 0)}{P(Y_i | W_i = 1)} = \frac{2}{\sigma^2} Y_i, \quad i = 1, \dots, n,$$

and subsequently we use ℓ_i to represent its realization.

Introduce the hard decision function $\theta(y)$ as $\theta(y) = 0$ if $y \geq 0$ and 1 otherwise. For \underline{Y} , denote the vector of $[\theta(Y_1), \dots, \theta(Y_n)] \in \mathbb{F}_2^n$ by $\theta(\underline{Y})$ for simplicity. As an exercise of Bayes' rule, we have

$$\begin{aligned} \text{if } y_i \geq 0, P(W_i = 1 | Y_i = y_i) &= \frac{1}{1 + \exp(\ell_i)}, \text{ and} \\ \text{if } y_i < 0, P(W_i = 0 | Y_i = y_i) &= \frac{1}{1 + \exp(-\ell_i)}. \end{aligned}$$

So the conditional probability that a transmitted bit is flipped by the hard decision of the channel output is given by

$$P(\theta(Y_i) \neq W_i | Y_i = y_i) = \frac{1}{1 + \exp(|\ell_i|)}. \quad (1)$$

B. GRAND

In a nutshell, the idea of GRAND is to find some $\underline{E} \in \mathbb{F}_2^n$, called an error pattern, so that by flipping the hard decisions $\theta(\underline{Y})$ according to \underline{E} , the result $\theta(\underline{Y}) \oplus \underline{E}$ is a codeword in \mathcal{C} .

In order to find such an error pattern that results in a codeword, a sequence of candidate error patterns are sequentially queried, until a codeword is reached, or until all candidate error patterns are exhausted. For ease of exposition, here we let the set of candidate error patterns be $\mathcal{E} = \{\underline{e}(1), \underline{e}(2), \dots, \underline{e}(2^n)\} = \mathbb{F}_2^n$, in which each $\underline{e}(t)$ is a different vector in \mathbb{F}_2^n . Hence an error pattern \underline{E} can surely be found, sooner or later. At the end of this subsection we will discuss how, in practice, we truncate the querying process due to complexity constraints.

The driving engine of a GRAND scheme is an ordering policy π , which is a permutation of $\{1, 2, \dots, 2^n\}$, depending upon \underline{y} . A GRAND scheme thus sequentially checks whether the criterion $\theta(\underline{y}) \oplus \underline{e}(\pi(t)) \in \mathcal{C}$ is satisfied, for t from 1 to 2^n .

It stops as soon as it finds the first t satisfying the criterion and declares the found codeword as the decoding result.

We generally determine the ordering policy π using a reliability vector $\gamma(\underline{y}) = [\gamma_1(\underline{y}), \dots, \gamma_n(\underline{y})]$, in such a way that, if any two candidate error patterns, say $\underline{e}(t)$ and $\underline{e}(t')$, satisfy

$$\sum_{i=1, \dots, n; e_i(t)=1} \gamma_i(\underline{y}) \leq \sum_{i=1, \dots, n; e_i(t')=1} \gamma_i(\underline{y}), \quad (2)$$

then $\pi(t) \leq \pi(t')$. That is, candidate error patterns are sorted in ascending order in terms of their accumulated reliability values.

The description using the reliability metric γ in the previous paragraph has led to a unified treatment of GRAND schemes [6]. Different choices of γ correspond to different GRAND schemes; for example, the original GRAND [1] has $\gamma_i(\underline{y}) = 1$, SGRAND [4] (which is equivalent to ML decoding) has $\gamma_i(\underline{y}) = |\ell_i|$ (see also [7, Ch. 10]), ORBGRAND [3] has $\gamma_i(\underline{y}) = r_i$ where r_i is the rank of $|\ell_i|$ among $\{|\ell_1|, \dots, |\ell_n|\}$, in ascending order, and so on.

We say that an ordering policy π is of ORB-type, if it only depends upon $\underline{r} = [r_1, \dots, r_n]$, instead of the exact values of \underline{y} . ORB-type ordering policies inherit the key advantage of ORBGRAND, namely its amenability to hardware implementation: once the relationship of ranking among the magnitudes of LLRs is obtained and the channel output symbols are rearranged accordingly, the queries can then be sequentially conducted using a pre-generated sequence, which can be efficiently implemented in a recursive fashion [3] [11]. Our aim is to design a new ORB-type ordering policy so as to approach the performance of SGRAND, which, unfortunately, is not of ORB-type.

III. SEARCH PROBLEM AND RESHUFFLED QUERYING

In order to conduct a tractable analysis of the querying process in GRAND schemes, we turn to a search problem. Compared with decoding, the key simplification is that the search problem ignores the case where a querying process stops with a codeword different from the actually transmitted one. Our analysis for such an idealized situation sheds key insight into how an ordering policy should be designed.

A. Search Problem Formulation

We formulate the search problem as follows:

- For the channel model described in Section II-A, let a codeword \underline{W} be randomly chosen, so as to produce the BPSK transmitted vector \underline{X} and the noise-corrupted channel output vector \underline{Y} , which is observed by the searcher.
- The searcher adopts an ordering policy π , which is a permutation of $\{1, 2, \dots, 2^n\}$, depending upon \underline{Y} .
- The searcher sequentially conducts queries by computing $\theta(\underline{Y}) \oplus \underline{e}(\pi(t))$, for t from 1 to 2^n . If for some t , $\theta(\underline{Y}) \oplus \underline{e}(\pi(t)) = \underline{W}$ is satisfied, a genie immediately informs the searcher that the chosen codeword \underline{W} is found; otherwise, the searcher turns to the next t .

Here lies the key difference between the search problem and GRAND: the searcher does not care about whether any other codeword is encountered; — instead, the searcher only cares about encountering \underline{W} during querying.

- For this idealized search problem, it is clear that there exists exactly one value of t such that $\theta(\underline{Y}) \oplus \underline{e}(\pi(t)) = \underline{W}$ is satisfied, and the searcher stops once this t is found.

B. Analysis of Search Problem

For the search problem, we are interested in the expected number of queries before the searcher stops. We emphasize that the expectation is with respect to the joint probability distribution of $(\underline{W}, \underline{Y})$, and hence it is an ensemble average result, given as follows.

Theorem 1. *For the search problem in Section III-A, the expected number of queries before the searcher stops is given by*

$$Q = \sum_{t=1}^{2^n} t \mathbf{E}_{\underline{Y}}[S_t], \quad (3)$$

where conditioned upon $\underline{Y} = \underline{y}$, S_t is realized as

$$s_t = P_{\underline{W}|\underline{Y}}(\theta(\underline{Y}) \oplus \underline{e}(\pi(t)) | \underline{Y} = \underline{y}), \quad (4)$$

in which the subscript $\underline{W}|\underline{Y}$ emphasizes that s_t is the posterior probability that the codeword is $\theta(\underline{Y}) \oplus \underline{e}(\pi(t))$ when the searcher observes \underline{Y} .

Proof: According to the search problem described in Section III-A, the expected number of queries Q can be calculated as

$$\begin{aligned} & \sum_{\underline{w} \in \mathcal{C}} P(\underline{W} = \underline{w}) \int_{\mathbb{R}^n} p_{\underline{Y}|\underline{W}}(\underline{y}|\underline{w}) \cdot \sum_{t=1}^{2^n} t \cdot \mathbf{1}(\theta(\underline{y}) \oplus \underline{w} = \underline{e}(\pi(t))) d\underline{y} \\ &= \int_{\mathbb{R}^n} p(\underline{Y} = \underline{y}) \sum_{\underline{w} \in \mathcal{C}} P_{\underline{W}|\underline{Y}}(\underline{w}|\underline{y}) \sum_{t=1}^{2^n} t \cdot \mathbf{1}(\theta(\underline{y}) \oplus \underline{w} = \underline{e}(\pi(t))) d\underline{y} \\ &= \int_{\mathbb{R}^n} p(\underline{Y} = \underline{y}) \sum_{t=1}^{2^n} t \cdot P_{\underline{W}|\underline{Y}}(\theta(\underline{y}) \oplus \underline{e}(\pi(t)) | \underline{y}) d\underline{y} \\ &= \sum_{t=1}^{2^n} t \cdot \int_{\mathbb{R}^n} p(\underline{Y} = \underline{y}) P_{\underline{W}|\underline{Y}}(\theta(\underline{y}) \oplus \underline{e}(\pi(t)) | \underline{y}) d\underline{y} \\ &= \sum_{t=1}^{2^n} t \mathbf{E}_{\underline{Y}}[S_t], \end{aligned}$$

where the second equality is due to that for given \underline{y} and $\underline{e}(\pi(t))$, there is exactly one \underline{w} , namely $\theta(\underline{y}) \oplus \underline{e}(\pi(t))$, for which the indicator function is one. ■

According to Theorem 1, the key to assessing a search problem is S_t , $t = 1, \dots, 2^n$. When the elements of \underline{W} , $[W_1, \dots, W_n]$, are independent and identically distributed (i.i.d.), S_t can be evaluated using the following result.

Lemma 1. *When $[W_1, \dots, W_n]$ are i.i.d., given \underline{y} , the posterior probability s_t in (4) is given by*

$$s_t = \prod_{i:e_i(\pi(t))=1} \frac{1}{1 + \exp(|\ell_i|)} \prod_{i:e_i(\pi(t))=0} \frac{\exp(|\ell_i|)}{1 + \exp(|\ell_i|)}. \quad (5)$$

Proof: Due to the i.i.d. property of \underline{W} and the memoryless property of the AWGN channel, an exercise of Bayes' rule yields

$$\begin{aligned} s_t &= P_{\underline{W}|\underline{Y}}(\theta(\underline{Y}) \oplus \underline{e}(\pi(t)) | \underline{Y} = \underline{y}) \\ &= \prod_{i=1}^n P_{W_i|Y_i}(\theta(Y_i) \oplus e_i(\pi(t)) | Y_i = y_i). \end{aligned} \quad (6)$$

Inspect the product. For those factors whose subscripts i 's satisfy $e_i(\pi(t)) = 1$, $\theta(Y_i)$'s are flipped, and according to (1), we have their product as $\prod_{i:e_i(\pi(t))=1} \frac{1}{1 + \exp(|\ell_i|)}$; on the other hand, for those factors whose subscripts i 's satisfy $e_i(\pi(t)) = 0$, $\theta(Y_i)$'s are not flipped, and according to (1), we have their product as $\prod_{i:e_i(\pi(t))=0} \frac{\exp(|\ell_i|)}{1 + \exp(|\ell_i|)}$. Multiplying these two parts together leads to (5). ■

From the expression of Q in (3), we have that ordering policies minimizing the expected number of queries should render the sequence of $\{\mathbf{E}_{\underline{Y}}[S_t]\}_{t=1, \dots, 2^n}$ monotonically non-increasing.

Corollary 1. *For the search problem, ordering policies that minimize Q satisfy $\mathbf{E}_{\underline{Y}}[S_1] \geq \mathbf{E}_{\underline{Y}}[S_2] \geq \dots \geq \mathbf{E}_{\underline{Y}}[S_{2^n}]$.*

Proof: This immediately follows from the expression (3). ■

The expectation sequence $\{\mathbf{E}_{\underline{Y}}[S_t]\}_{t=1, \dots, 2^n}$ being monotone is a condition weaker than the sequence $\{s_t\}_{t=1, \dots, 2^n}$ itself being monotone for each \underline{y} . In fact, this latter condition corresponds to the ordering policy of SGRAND, as shown by the following result.

Proposition 1. *The ordering policy π that for each \underline{y} rearranges queries so as to render $s_1 \geq s_2 \geq \dots \geq s_{2^n}$ is identical to the ordering policy of SGRAND generated using $\gamma_i(\underline{y}) = |\ell_i|$ in Section II-B.*

Proof: Consider any two candidate error patterns, $\underline{e}(t)$ and $\underline{e}(t')$, satisfying $s_t \geq s_{t'}$. Applying Lemma 1, we have

$$\begin{aligned} & \prod_{i:e_i(\pi(t))=1} \frac{1}{1 + \exp(|\ell_i|)} \prod_{i:e_i(\pi(t))=0} \frac{\exp(|\ell_i|)}{1 + \exp(|\ell_i|)} \\ & \geq \prod_{i:e_i(\pi(t'))=1} \frac{1}{1 + \exp(|\ell_i|)} \prod_{i:e_i(\pi(t'))=0} \frac{\exp(|\ell_i|)}{1 + \exp(|\ell_i|)}. \end{aligned}$$

Noting that the denominators on both sides of the inequality are in fact identical, i.e., $\prod_{i=1}^n (1 + \exp(|\ell_i|))$, we cancel them and obtain

$$\begin{aligned} & \prod_{i:e_i(\pi(t))=0} \exp(|\ell_i|) \geq \prod_{i:e_i(\pi(t'))=0} \exp(|\ell_i|), \\ \text{i.e., } & \sum_{i:e_i(\pi(t))=0} |\ell_i| \geq \sum_{i:e_i(\pi(t'))=0} |\ell_i|, \\ \text{i.e., } & \sum_{i:e_i(\pi(t))=1} |\ell_i| \leq \sum_{i:e_i(\pi(t'))=1} |\ell_i|, \end{aligned} \quad (7)$$

which is exactly the condition of $\pi(t) \leq \pi(t')$ in SGRAND, according to the description in Section II-B. ■

C. Reshuffled Querying

Motivated by Proposition 1 and Corollary 1, we desire an ORB-type ordering policy that approximates the ordering policy of SGRAND in the sense that, although it is unlikely to have $\{s_t\}_{t=1,\dots,2^n}$ be monotonically non-increasing for each y , the overall effect, i.e., the ensemble average sequence $\{\mathbf{E}_Y[\hat{S}_t]\}_{t=1,\dots,2^n}$ is monotonically non-increasing. So we propose the following simple approach:

- We begin with an existing ORB-type ordering policy, denoted π_{base} , for example, the original ORBGRAND [3] or CDF-ORBGRAND [3] [6], to obtain its corresponding expectation sequence $\{\mathbf{E}_Y[S_t]\}_{t=1,\dots,2^n}$.
- We then sort $\{\mathbf{E}_Y[S_t]\}_{t=1,\dots,2^n}$ so that they are in descending order. Denote the permutation achieving this goal by $\tilde{\pi}$.
- We then use the concatenation of π_{base} and $\tilde{\pi}$, $\pi = \tilde{\pi} \circ \pi_{\text{base}}$, as the ordering policy for decoding.

We call the thus obtained decoder RS-ORBGRAND, where the prefix “RS” implies that it is obtained by reshuffling the ordering policy of a base ORBGRAND scheme. Since both π_{base} and $\tilde{\pi}$ are of ORB-type, RS-ORBGRAND is of ORB-type. In particular, we emphasize that, conditioned upon π_{base} , $\tilde{\pi}$ is in fact a fixed permutation, even independent of the ranking \mathcal{L} , let alone y . For obtaining $\tilde{\pi}$, only the expectation sequence $\{\mathbf{E}_Y[S_t]\}_{t=1,\dots,2^n}$ is needed, which can be evaluated offline, without requiring the exact realizations y .

RS-ORBGRAND uses the reshuffling step $\tilde{\pi}$ to render the queries to satisfy Corollary 1. For each particular y , this reshuffling may not always render the queries to satisfy $s_1 \geq s_2 \geq \dots \geq s_{2^n}$. So there still exists a gap between RS-ORBGRAND and SGRAND. The following result sheds some further insight into this gap.

Proposition 2. *For any ordering policy π , the expectation of its excess number of queries compared with SGRAND is given by*

$$\Delta Q = \mathbf{E}_Y \left[\sum_{1 \leq i < j \leq 2^n : S_i < S_j} (S_j - S_i) \right]. \quad (8)$$

Proof: Denote by \hat{S}_t , $t = 1, \dots, 2^n$, the random variables characterized by (4) induced by SGRAND. In fact, $\{\hat{S}_t\}_{t=1,\dots,2^n}$ is simply a rearrangement of $\{S_t\}_{t=1,\dots,2^n}$ satisfying $\hat{S}_1 \geq \dots \geq \hat{S}_{2^n}$. Applying the expression of Q in Theorem 1, we have that the difference between the expected numbers of queries with the considered π and SGRAND, i.e., the expectation of the excess number of queries, is given by

$$\Delta Q = \mathbf{E}_Y \left[\sum_{t=1}^{2^n} t(S_t - \hat{S}_t) \right]. \quad (9)$$

To prove the equivalence between (8) and (9), we compare their coefficients of S_t . Denote by ρ_t the position of S_t in $\{\hat{S}_t\}_{t=1,\dots,2^n}$. In (9), the coefficient of S_t is $t - \rho_t$. On the other hand, in (8), the coefficient of S_t can be rewritten as $\sum_{i:i < t} \mathbf{1}(S_i < S_t) - \sum_{j:j > t} \mathbf{1}(S_j > S_t)$.

We now take a closer look at how we rearrange $\{S_t\}_{t=1,\dots,2^n}$ to obtain $\{\hat{S}_t\}_{t=1,\dots,2^n}$. Starting with S_t , since there are $\sum_{j:j > t} \mathbf{1}(S_j > S_t)$ elements with indices larger than t and values also larger than S_t , we need to move them to “the left of” S_t , and since there are $\sum_{i:i < t} \mathbf{1}(S_i < S_t)$ elements with indices smaller than t and values also smaller than S_t , we need to move them to “the right of” S_t . This way, S_t is moved to the position ρ_t as \hat{S}_{ρ_t} . Therefore, we have

$$t + \sum_{j:j > t} \mathbf{1}(S_j > S_t) - \sum_{i:i < t} \mathbf{1}(S_i < S_t) = \rho_t, \\ \text{i.e.,} \quad t - \rho_t = \sum_{i:i < t} \mathbf{1}(S_i < S_t) - \sum_{j:j > t} \mathbf{1}(S_j > S_t),$$

meaning that the coefficients of S_t in (8) and (9) are identical, for each $t = 1, \dots, 2^n$. ■

We can rewrite ΔQ in (8) as

$$\Delta Q = \sum_{i=1}^{2^n} \sum_{j=i+1}^{2^n} \mathbf{E}_Y[(S_j - S_i) \cdot \mathbf{1}(S_j > S_i)], \quad (10)$$

and construct a $2^n \times 2^n$ matrix \mathcal{R} with its (j, i) -th element as:

$$\mathcal{R}_{j,i} = \begin{cases} \mathbf{E}_Y[(S_j - S_i) \cdot \mathbf{1}(S_j > S_i)] & i < j \\ 0 & i = j \\ \mathbf{E}_Y[(S_i - S_j) \cdot \mathbf{1}(S_i > S_j)] & i > j \end{cases} \quad (11)$$

We see that ΔQ is the sum of the lower-triangular elements of \mathcal{R} . The reshuffling step $\tilde{\pi}$ in RS-ORBGRAND basically keeps swapping rows and columns of \mathcal{R} , so that if for any $i < j$, $\mathcal{R}_{j,i} > \mathcal{R}_{i,j}$, these two elements are interchanged. Using \mathcal{R} , we can visually illustrate the effect of the reshuffling step $\tilde{\pi}$. In Figure 1, we display the first 200×200 rows and columns of \mathcal{R} , before and after applying the reshuffling step $\tilde{\pi}$ to the base π_{base} , for which we choose to use CDF-ORBGRAND. Here for better visual effect, we plot in gray scale the normalized values $\mathcal{R}_{j,i}/(\mathcal{R}_{j,i} + \mathcal{R}_{i,j})$. It can be seen that after reshuffling, all the lower-triangular elements of \mathcal{R} become no greater than 0.5, i.e., all the relatively large lower-triangular elements are interchanged with their upper-triangular counterparts.




Fig. 1. Illustration of \mathcal{R} before and after applying the reshuffling step.

In practice, since 2^n is typically an exceedingly large value, due to complexity constraints, we cannot afford the worst case. In implementation, we introduce two additional parameters, $T \leq T_1$ no greater than 2^n . In RS-ORBGRAND,

TABLE I
AVERAGE NUMBER OF QUERIES FOR BCH(127, 113).

| | 4dB | 5dB | 6dB | 7dB |
|------------------------|-------|-------|-------|-------|
| ORBGRAND [3] | 790.8 | 83.89 | 7.072 | 1.479 |
| CDF-ORBGRAND [3] | 727.9 | 67.44 | 5.476 | 1.478 |
| 3 Line-ORBGRAND [3] | 730.0 | 62.34 | 4.732 | 1.445 |
| RS-ORBGRAND (proposed) | 715.6 | 60.63 | 4.445 | 1.350 |
| SGRAND [2] | 666.5 | 52.99 | 3.932 | 1.328 |

we apply the reshuffling step to only the first T_1 elements of $\{\mathbf{E}_Y[S_t]\}_{t=1,\dots,2^n}$. When decoding, we truncate the querying process if none of $\mathbf{e}(\pi(t))$, $t = 1, \dots, T$, yields a codeword; — when truncation occurs, no codeword is found and a decoding failure is declared. As will be seen in the next section, choosing a value of T_1 sufficiently larger than T is crucial for achieving a good performance. Noting that the reshuffling step is conducted offline, we can usually tolerate to use a sufficiently large value of T_1 , given the abundant computing resources available to date.

IV. PERFORMANCE EVALUATION

In this section, we apply RS-ORBGRAND proposed in the previous section to conduct some numerical studies. We consider two codes: BCH(127, 113) and polar(128, 114) with 10 cyclic redundancy check (CRC) bits. Unless specified otherwise, we use CDF-ORBGRAND [3] [6] with $T_1 = 5 \times 10^4$ to generate the ordering policy, and when decoding, we permit at most $T = 10^4$ queries and terminate the querying process if no codeword is still not found by then.

Figure 2 compares the BLER for BCH(127, 113) using several ORB-type decoders as well as SGRAND and ML decoding lower bound. The ML decoding lower bound is obtained by first running SGRAND with $T = 10^5$, and then supposing that even if no codeword has been found upon truncation, the decoding result would still be correct. It can be seen that only RS-ORBGRAND achieves performance close to SGRAND as SNR increases. When BLER is 10^{-6} , the gain of RS-ORBGRAND compared with existing ORB-type decoders is at least 0.3dB, and the gap from the ML decoding lower bound is only 0.1dB. In Table I we display the average number of queries for the decoders studied in Figure 2, and we can observe that RS-ORBGRAND requires the least querying cost among ORB-type decoders. We remark that although SGRAND incurs the least average number of queries (as also suggested by the analysis in Section III-C), its queries needs to be generated online based on the exact values of LLRs and hence its querying complexity is by far the highest.

Figure 3 compares the BLER for CRC-aided polar(128, 114) using different decoders. Here for RS-ORBGRAND, we further study several different values of T_1 , from 5×10^4 down to 10^4 . It is clearly shown that using a sufficiently large set of candidate error patterns is crucial. Again, we note that T_1 candidate error patterns are used offline for obtaining the reshuffling step, but when decoding, all the decoders can




Fig. 2. BLER for BCH(127, 113).




Fig. 3. BLER for CRC-aided Polar(128, 114).

conduct at most $T = 10^4$ queries. By using a large value of T_1 , we can identify more relatively large $\mathbf{E}_Y[S_t]$'s and reshuffle them, so as to have them queried earlier in RS-ORBGRAND.

V. CONCLUSION

RS-ORBGRAND, as an improvement of ORBGRAND, is proposed and studied. Its basic idea is to reshuffle queries with the goal of approximating SGRAND. This improved scheme still inherits the key advantage of ORBGRAND, namely its dependency upon the ranking of the magnitudes of LLRs only. Numerical simulations show that RS-ORBGRAND achieves noticeable gain compared with existing ORB-type decoders, and is very close to ML decoder even in the regime of low BLER.

REFERENCES

- [1] K. R. Duffy, J. Li, and M. Médard, "Capacity-achieving guessing random additive noise decoding," *IEEE Transactions on Information Theory*, vol. 65, no. 7, pp. 4023–4040, 2019.
- [2] K. R. Duffy, M. Médard, and W. An, "Guessing random additive noise decoding with symbol reliability information (SRGRAND)," *IEEE Transactions on Communications*, vol. 70, no. 1, pp. 3–18, 2021.
- [3] K. R. Duffy, W. An, and M. Médard, "Ordered reliability bits guessing random additive noise decoding," *IEEE Transactions on Signal Processing*, vol. 70, pp. 4528–4542, 2022.
- [4] A. Solomon, K. R. Duffy, and M. Médard, "Soft maximum likelihood decoding using GRAND," in *IEEE International Conference on Communications (ICC)*, 2020, pp. 1–6.
- [5] C. Yue, V. Miloslavskaya, M. Shirvanimoghaddam, B. Vucetic, and Y. Li, "Efficient decoders for short block length codes in 6G URLLC," *IEEE Communications Magazine*, vol. 61, no. 4, pp. 84–90, 2023.
- [6] M. Liu, Y. Wei, Z. Chen, and W. Zhang, "ORBGRAND is almost capacity-achieving," *IEEE Transactions on Information Theory*, vol. 69, no. 5, pp. 2830–2840, 2022.
- [7] S. Lin and D. J. Costello, *Error Control Coding: Fundamentals and Applications*. Englewood Cliffs, NJ, USA: Prentice Hall, 2004.
- [8] K. Galligan, M. Médard, and K. R. Duffy, "Block turbo decoding with ORBGRAND," in *57th Annual Conference on Information Sciences and Systems (CISS)*, 2023, pp. 1–6.
- [9] W. An, M. Médard, and K. R. Duffy, "Soft decoding without soft demapping with ORBGRAND," in *2023 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2023, pp. 1080–1084.
- [10] S. M. Abbas, T. Tonnellier, F. Ercan, M. Jalaleddine, and W. J. Gross, "High-throughput and energy-efficient VLSI architecture for ordered reliability bits GRAND," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 30, no. 6, pp. 681–693, 2022.
- [11] C. Condo, V. Bioglio, and I. Land, "High-performance low-complexity error pattern generation for ORBGRAND decoding," in *IEEE Globecom Workshops (GC Wkshps)*, 2021, pp. 1–6.
- [12] S. M. Abbas, M. Jalaleddine, and W. J. Gross, "List-GRAND: A practical way to achieve maximum likelihood decoding," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 31, no. 1, pp. 43–54, 2022.
- [13] C. Condo, "A fixed latency ORBGRAND decoder architecture with LUT-aided error-pattern scheduling," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 69, no. 5, pp. 2203–2211, 2022.